Essays on Information Conveyance

Daniel Venture Kwiatkowski

Charlottesville, VA

M.A. Economics, University of Virginia, 2020

B.A. Economics, University of Virginia, 2019

A Dissertation presented to the Graduate Faculty

of the University of Virginia in Candidacy for the Degree of

Doctor of Philosophy

Department of Economics

University of Virginia

May 2025

Charlie Holt

Simon Anderson

Maxim Engers

Noah Myung

To Julia

Acknowledgements

I am grateful to Charlie Holt, Simon Anderson, Maxim Engers, and Noah Myung for their guidance. I also want to thank my wife Julia Kwiatkowski as well as Joe Anderson, Diego Briones, and Sasha Ruby. I'm grateful for feedback and encouragement from Marc Santugini and from the members of the experimental reading group at UVA, especially Yen Ling Tan, Zitian Wang, and Po-Hsuan Lin. This work was funded in part by the Bankard Fund for Political Economy Pre-Doctoral Fellowship.

Table of Contents

Introduction

Can prices convey information about product quality? Empirically, higher-quality products tend to be more expensive, and consumers often use prices as signals of quality. Despite consumers' apparent belief that prices communicate otherwise inaccessible quality information, this is difficult to justify intuitively. Price is a single decision variable, yet it must serve multiple purposes—it must serve the price-setters profit incentives while also, potentially, signaling quality. Informally, expecting price to play both roles is like expecting a single variable to solve two unrelated equations.

In my work, I formalize this idea that price cannot "do two jobs at once." Prices are set by sellers who have a single overarching objective: to maximize profits. However, profit maximization consists of two competing components: (a) maximizing profits given buyers' existing beliefs and (b) influencing those beliefs to induce highest willingness to pay. Ideally, a seller would like to achieve both objectives simultaneously, but a single pricing strategy cannot do so. A price that credibly signals high quality may not maximize profits when the product is low quality.

Consider a seller who sometimes offers high-quality products and sometimes low-quality ones. If they price high for good products and low for bad ones, they train buyers to associate high prices with high quality, allowing them to maximize profits when selling high-quality goods.[1] However, once buyers form this belief, the seller has an incentive to charge a high price even when selling low-quality products. If no external mechanisms (such as certifications, reviews, or repeat business) prevent deception, the seller's optimal strategy is to exploit buyers' trust—undermining the credibility of the price signal.

[1]This is not necessarily a process that happens over time, it is simply the result of an equilibrium where beliefs are consistent.

As a result, perfect information transmission through prices is generally impossible. However, this does not mean that prices convey no information at all. Rather, the amount of information that can be credibly conveyed is limited by the incentive-compatibility constraints that prevent sellers from finding deception too profitable. Across several models, these constraints place an upper bound on how much information can be transmitted in equilibrium.

All three essays relate to this core idea. In the first, I run experiments to see the Grossman-Stiglitz Paradox in the lab. The upshot of the paradox is that prices can be informative, but not *too* informative: "the impossibility of informationally efficient markets". Prices can be *just informative enough* that the return to arbitrage perfectly offsets its cost. This same intuition applies to the buyer-seller version of the Grossman-Stiglitz paradox that I implement (originally from Bester and Ritzberger): prices can be just informative enough that the return to investigating product quality offsets its cost. I find that subjects arrive at this "informative but not too informative" outcome. With a few behavioral caveats, they are very close to what the G-S Paradox predicts.

But this idea is broader than the Grossman and Stiglitz (1980) model; it also applies to models where arbitrage is *not* the mechanism for informativeness. The second essay investigates a model in which buyers have no way to become informed. Prices can still reflect quality information, however, since high-quality products are more costly for sellers to produce. Again, prices can be informative, but not so informative that sellers always try to cheat buyers. Subjects arrive at somewhat informative pricing in the lab, and I examine how this varies as competition between sellers increases.

The third essay, analyzing jury voting games, is a different sort of model. But, the intuition is similar: actions (in this case, voting choices rather than prices) convey some information, but they cannot convey too much information. When votes are perfectly informative, it's because people vote according to their types. But this doesn't play nicely with, say, a unanimous voting rule. People have just one lever (their voting decision) both to convey information to others, and to get the correct aggregate outcome. That one lever cannot do both things. It is similar to how sellers have one lever (the price) and they want to use it both to convey their high quality to sellers but also to maximize their payoff. If a buyer sees a high price, they know the seller's need to maximize their payoff could have pushed them to set a high price even if it doesn't accurately signal their type. In the same way, when you see a voter vote yes, you know their incentive to influence the overall unanimous voting outcome could have pushed them to vote yes even if this doesn't accurately convey their private information.

Chapter 1

Can Price Inform Quality when Verification is Costly?

Abstract

A product's price will reflect its quality, if that quality is known to consumers. But many experiments have shown that consumers believe high prices signal high quality, even in situations where the quality is *not* generally known. It's unclear if consumers are right or wrong to expect this. I examine informational efficiency in the lab, using a market where verifying product quality is costly. Theoretically, prices could convey information according to Grossman and Stiglitz (1980), but might also convey no information. I find that prices endogenously convey about as much information as theoretically possible, even when information is quite costly. Behavioral bias, which I examine using quantal response, can make prices slightly more or less informative than theory predicts.

1 Introduction

Prices are often correlated with product quality. Ceteris parabus, a first-class airline ticket is more expensive than economy class, a house with a beautiful view is more expensive than one without, and an acclaimed bottle of French wine is more expensive than a boxed wine grown in Houston. Does this mean if I encounter two otherwise indistinguishable wines, one more expensive than the other, I can surmise that the more expensive wine will be better than the cheaper wine? Not necessarily. It may be that prices reflect quality only because quality information is already obvious to consumers, so that if quality is *not* apparent, there is no reason for the better product to be more expensive.

In general, it is difficult to say whether prices are doing informational work: successfully conveying quality information to otherwise uninformed consumers. But in many cases, consumers believe that prices can do informational work. In the lab, it has been shown that buyers believe prices convey quality information even when the quality is not obvious otherwise (Leavitt 1954, McConnell 1968, Olson 1977). These experiments are focused on understanding the buyers' beliefs, and involve experimenters varying the price of a product as a treatment. The experiments show that buyers believe prices are somewhat informative of quality, but it's unclear if buyers are correct to believe that prices would be informative when quality is not obvious, or if their beliefs are an artifact from real-world markets in which quality is generally known (or even a reflection of what situations they expect experimenters to present them with).

This paper examines whether prices convey quality information endogenously in a laboratory experiment where quality is not freely observable. There are two reasons why it might be valuable to see informative pricing arise endogenously in the lab. The first is to understand the mechanism by which prices become informative. Theoretically, there are many models that lead to informative pricing and in real-world markets it is generally not clear which channels, if any, are operative.

The other reason to examine endogenous price information in the lab is to see how behavioral biases may impact how much information prices convey. Noise in decision-making, risk-aversion, and learning may make prices more or less informative than theoretical models predict. In the lab, behavioral biases can be seen clearly, while in real-world markets their effects may be confounded with variation from other sources such as heterogeneity in consumers' values, producers' costs, and product characteristics.

I examine the informational content of prices through the channel studied in Grossman and Stiglitz (1980), and I follow Bester and Ritzberger (2001) in adapting this intuition to a buyer-seller framework. Here, prices convey some information in equilibrium because consumers can exert effort to verify product quality. I find that this channel is operative in the lab. In every experimental session, subjects reach the separating equilibrium in which prices are informative. I also find that the gradual responsiveness of subject behavior, relative to the sharp discontinuities in theoretical best responses, can make prices either more or less informative than theory predicts.

These results demonstrate the efficient markets hypothesis as stated by Grossman and Stiglitz. The prices observed in the lab are about as informative as they can be theoretically.

An individual subject will update their beliefs about quality after seeing the price and will then be indifferent between either verifying the quality or buying the product without verifying.

But there are also behavioral barriers to efficiency. Subjects are very responsive to the cost of verifying product quality. When verification is cheap, buyers verify more frequently than necessary; ignoring that, since others are also verifying, the price already reflects information about quality. Instead, buyers do not trust the market enough. When information is expensive, the opposite happens. Buyers get information rarely, and trust the market too much, failing to realize that, since others are also getting less information, the price is a less reliable guide to quality.

2 Background

Formulation of the efficient markets hypothesis by Samuelson (1965) and Fama (1970) generally assumed that information was freely available to investors who could then bid the price up or down through buying or selling decisions, until no further arbitrage was profitable. Grossman and Stiglitz (1980) pointed out that, if information is freely observable, then it is immaterial whether prices convey information. Prices are doing meaningful informational work only when they are conveying information to otherwise uninformed buyers. Grossman and Stiglitz thus assumed arbitrage was costly and reformulated the efficient markets hypothesis so that the return to arbitrage perfectly offset its cost. Markets were then "efficient" if prices conveyed as much information as possible in equilibrium—just enough to make buyers indifferent to arbitrage opportunities. Informative pricing is sustained by buyers who incur costs in order to benefit from the deviation of current prices from expected returns.

Since Grossman and Stiglitz, many other models have appeared that sustain informative pricing. Some involve quality entering demand (through, for instance, a proportion of informed consumers, as in Bagwell and Riordan, 1991). Some involve quality entering through a firm's cost (e.g. Tirole, 1988, pp. 107-108). In 2001, Bester and Ritzberger made the Grossman and Stiglitz intuition more concordant with these models by adapting it for one buyer and one single-product monopolist. This gives the model the same separating/pooling dimension as Bagwell and Riordan but without a proportion of already informed consumers. Instead, arbitrage is captured by buyers who can exert costly effort to learn the true quality of the product.

I have chosen to use a discretized version of Bester and Ritzberger's model to test price informativeness in the lab. One reason for this is that the model is simple and easy for subjects to understand. Furthermore, people in real-world markets do expend effort to learn about the quality of products prior to buying. So, this is certainly one active informational channel, though it may not be the only one.

3 The Model

A seller is endowed with a single product of quality v, known only to themselves, and chooses a price p at which to sell the product. If the product is not sold, the seller gets no payoff—the product is useless to the seller. If the product does sell, the seller gets the price they set. Thus the seller's payoff is

$$\pi_{\text{seller}} = p\mathbb{1}\{B\}$$

where $\mathbb{1}\{B\}$ is an indicator for whether the buyer buys.

The buyer knows the prior distribution of v and observes the price set by the seller. The buyer then chooses one of three options. The buyer can buy the product, in which case they receive v–p, or they can walk away without buying, in which case they receive 0. The buyer can alternatively incur an effort cost of c to observe v. The cost c is not a transfer to the seller, it is simply a loss, representing the cost of time or other resources used. If the buyer pays c to observe v, they can then decide to buy and receive v–p–c or walk away and receive

$-c$. Overall, the buyer's payoff is

$$\pi_{\text{buyer}} = \mathbb{1}\{B\}(v-p) - \mathbb{1}\{I\}c$$

where $\mathbb{1}\{B\}$ is an indicator for whether the buyer buys and $\mathbb{1}\{I\}$ is an indicator for whether they pay for information.

To implement this game in the experiment, I restrict v to be either a high or low quality, and restrict the seller to two possible price choices, a high price and a low price. The timing of the game is as follows (the full extensive form is in the appendix). First, nature chooses the seller's quality:

$$v = \begin{cases} v_l \text{ w.p. } 1/2 \\ v_h \text{ w.p. } 1/2 \end{cases}$$

Then the seller chooses $p \in \{p_l, p_h\}$ where

$$p_l < v_l < \frac{v_l + v_h}{2} < p_h < v_h$$

After observing p, the buyer chooses whether to buy (B), leave (L), or get information (I). If the buyer chooses to get information, they observe v and then choose to buy or leave.

3.1 Equilibria

I will focus on the perfect Bayesian equilibria of the game. First, notice that once the buyer gets information, the next choice is straightforward. Once the cost of information is sunk, the buyer should buy whenever $v > p$. Next, note that the buyer should buy blindly whenever $p = p_l$. Since $p_l < v_l < v_h$, regardless of the quality of the product, it is worthwhile to buy. There is no benefit from acquiring information in this case. Taking these two things into account, we can rewrite the game tree (fig 1). The only difficult question for the buyer is

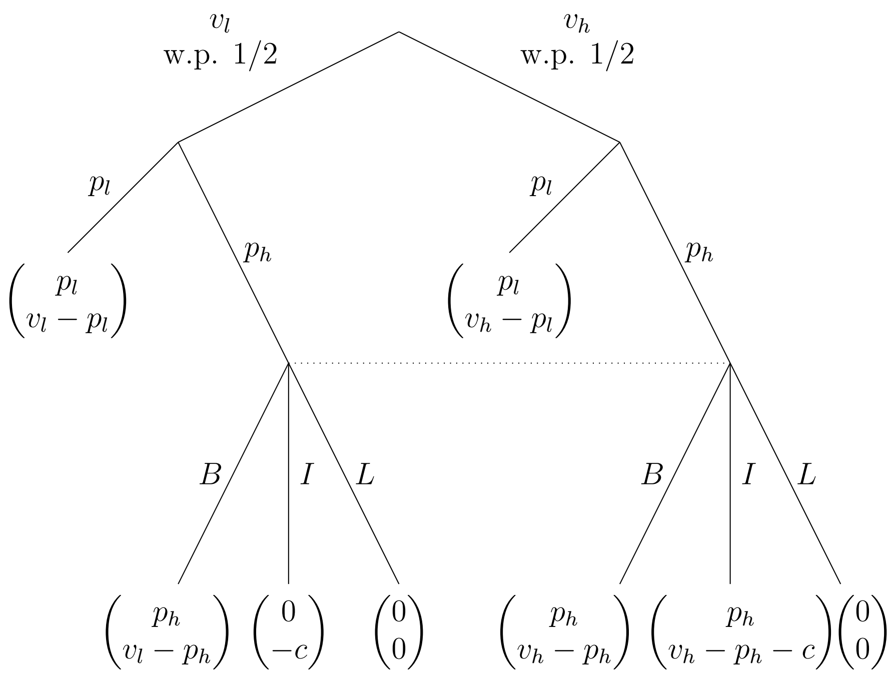

Figure 1: Should Buyers Buy Expensive Products?

Notes: The first decision is a move by nature that determines if the seller has high- or low-quality products. Then the seller chooses a price. If the seller chooses the low price, the buyer should buy. If the seller chooses the high price, the buyer must choose either buy (B), get information (I), or leave (L), depending on their posterior beliefs about the type of seller that would set a high price.

what to do when $p = p_h$.

If a buyer observes the high price, should they buy blindly, hoping it is a high-quality product? Should they walk away, assuming it isn't worth the risk? Or should they incur costly effort to check the quality? The answer depends on the buyer's beliefs. Suppose the buyer has updated beliefs $\mu_h \equiv \mathbb{P}(v = v_h | p = p_h)$ after seeing the high price. Then if the buyer buys blindly, they get

$$\mathbb{E}[v|p_h] – p_h = \mu_h v_h + (1 - \mu_h) v_l – p_h = v_l – p_h + (v_h - v_l)\mu_h$$

If the buyer gets information, they pay c but have a μ_h chance of getting $v_h – p_h$ if the product is high-quality, so their expected payoff is

$$\mu_h(v_h – p_h) - c$$

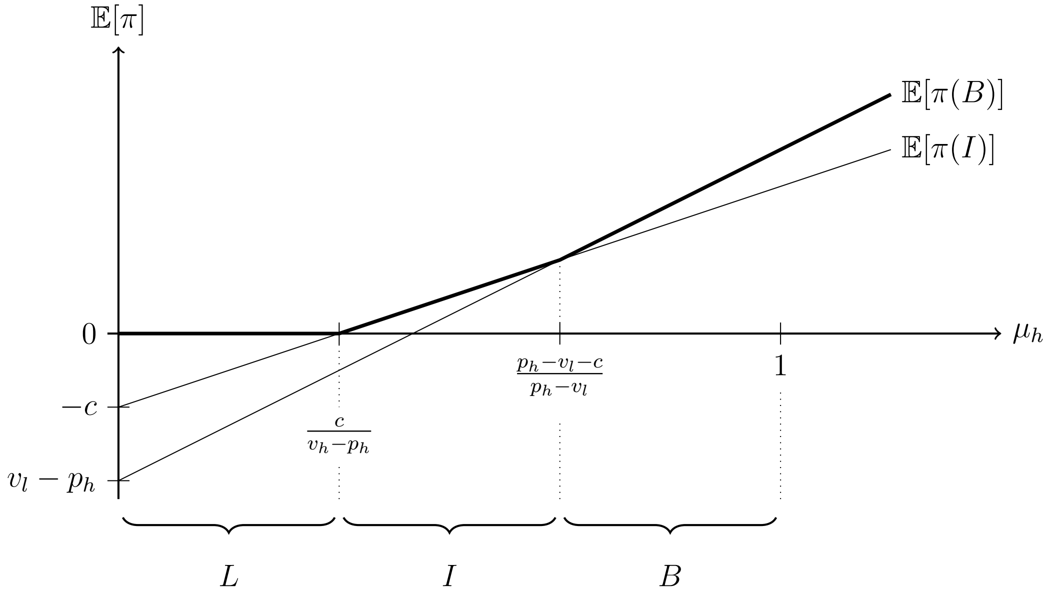

Figure 2: Optimal Buyer Choice for Different Beliefs

Notes: μ_h is the buyer's belief that a seller setting the high price has a high-quality product. If the buyer thinks high-priced products are very likely to be high-quality, they should buy (B). If they think high-priced products are unlikely to be high-quality, they should leave (L). For intermediate beliefs, they should choose to get information (I).

Figure 2 shows the optimal buyer best response to $p = p_h$ for different values of beliefs μ_h. If it is very likely that a high price indicates a high quality product, the buyer should just buy without checking. If it is more likely that the quality might be low, the buyer should check before buying. If it is almost certain that the product is low-quality, then it is not even worthwhile to check, and the buyer should just leave.

Suppose μ is low, so the buyer thinks a high price is quite likely to have come from a low-type seller. Then the buyer will leave without buying or getting information when $p = p_h$. If this is the case, sellers will have to set $p = p_l$ to make a sale, and this equilibrium is sustained by off-equilibrium-path beliefs: the buyer is justified in believing high prices may come from low-quality sellers because high prices never occur.[1] This is an uninformative equilibrium because the buyer learns nothing from the price—it conveys no information about the quality of the product.

[1]Note that this pooling equilibrium still satisfies the Cho-Kreps intuitive criterion because both high and low-type sellers would benefit from deviating to the high price if the buyer were to buy.

If instead μ_h is higher, so that the buyer thinks the high price is more likely to come from a high-type seller, then they will buy blindly or get information. If they always bought blindly, the low-type firm would have an incentive to fool the buyer by setting $p = p_h$ as well, and prices would no longer be informative. This is not an equilibrium because $p_h > (v_l + v_h)/2$, so if both types are setting $p = p_h$, it is not worthwhile to buy.

If the buyer always got information when seeing the high price, the low-type seller would have to set $p = p_l$ while the high-type seller could set $p = p_h$. But then the price would be perfectly informative, and paying for information would be a waste—the product quality is already obvious from the price. So perfect type dependence is also not an equilibrium.

Suppose μ_h is exactly at the threshold between buying blindly and getting information, so that the buyer is indifferent between the two. Here there is a partially separating equilibrium. The buyer mixes between the two strategies and the low-type seller makes this mixing plausible by occasionally trying to fool the buyer by setting $p = p_h$. Prices are as informative as possible in this equilibrium. The price conveys just enough information so that buyers are indifferent between trusting the market and doing their own research. The return to verification would just offset the cost, just as the return to arbitrage perfectly offsets its cost in the Grossman and Stiglitz model.

Thus, there are two types of equilibria, one where prices convey information and one where they do not.[2] While both are theoretically valid, I can use the experimental data to empirically test if subjects endogenously reach informative pricing or if they behave closer to the pooling equilibrium.

3.2 Comparative Statics

As well as examining equilibrium selection, I vary the cost of verification, c, to see how closely subjects track theoretical predictions. In the theory, as c decreases and information

[2]There is no equilibrium when μ_h is at the threshold where the buyer is indifferent between leaving and getting information: in this case, only high-quality products would be sold at high-prices, so the buyer would regret not buying them blindly. Nevertheless, there is not an even number of equilibria, since there is a continuum of pooling equilibria (one for each value of buyer beliefs).

is cheaper, prices become more and more informative until the price perfectly reflects quality when $c = 0$. As c increases and information becomes more expensive, prices become less informative until the informative equilibrium disappears.

However, this depends on sharp comparative statics characteristic of mixed Nash equilibria. As c increases (decreases), the low-type seller should try to fool the buyer more (less) to keep the buyer indifferent between buying blindly and getting information. The buyer should get information at the same rate regardless of the cost of information, since their behavior needs to keep the low-type seller indifferent in equilibrium.

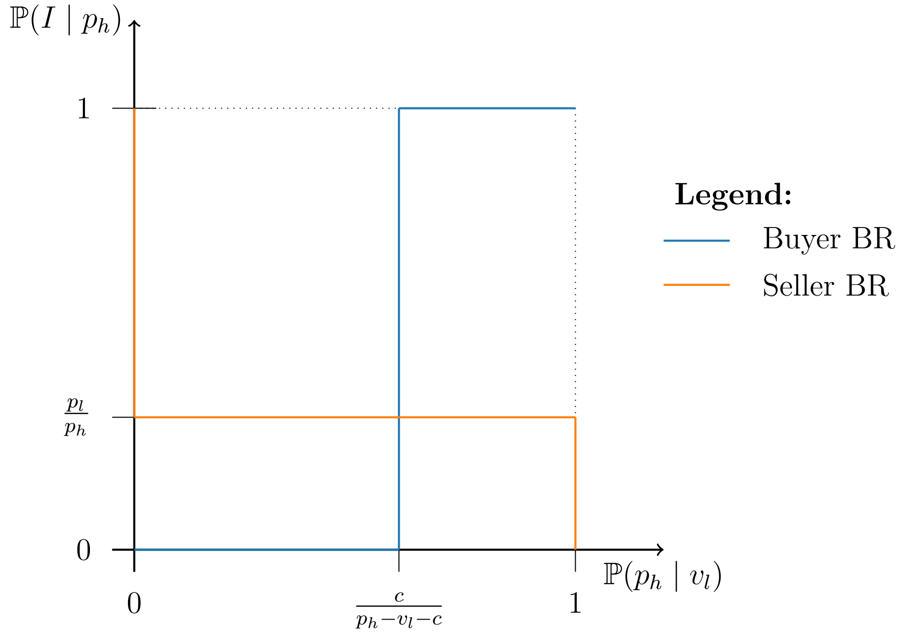

Figure 3: Cheating and Verifying is a Game of Chicken

Notes: This is the slice of the strategy space where the low-quality seller decides whether to "cheat" the buyer by setting the high price, and the buyer decided whether to verify the quality of a high-priced product. This is essentially a game of chicken, or cat-and-mouse. If the low-quality seller is frequently trying to cheat the buyer, the buyer should verify. But if the low-quality seller is being honest, the buyer should trust them and not exert effort to verify the quality. Conversely, if the buyer is trusting, the low-quality seller should cheat them, and if they are verifying, the low-quality seller should be honest. The only Nash equilibrium is in mixed strategies: the seller cheats just enough that the buyer is indifferent to getting information, and the buyer verifies just enough that the seller is indifferent. This equilibrium is unstable in the sense that a small deviation from one player should lead to a larger deviation in response, until players are far from the Nash equilibrium.

Figure 3 shows this result. This figure is just a slice of the actual action space: the slice where the buyer always buys when observing the low price, always either buys or verifies when observing the high price, and where the high-quality seller always sets the high price. These are the choices that should occur under the partially separating Nash equilibrium, and players usually follow these actions. This slice represents, essentially, a game of chicken (or cat-and-mouse) between the buyer and the low-quality seller. The buyer wants to buy given p_h if the low-quality seller sets p_l, and wants to verify given p_h if the low-quality seller sets p_h. The low-quality seller wants the opposite: they want to set p_l if the buyer is verifying and set p_h when the buyer is buying blindly.

But mixed Nash equilibria can be difficult for subjects to understand, and can sometimes require very astute subjects to achieve. Insofar as these Nash comparative statics are counterintuitive to subjects, empirical price informativeness may differ from theoretical predictions.

4 Experimental Design

Experiments were run in-person at the University of Virginia, with a sample of 74 undergraduate students. There were six sessions of 10-14 subjects each. In each session, subjects were randomly chosen to be a buyer or a seller, and this designation persisted throughout the session. Subjects then played the game for 16 rounds with one information cost, and another 16 rounds with a different information cost. At the beginning of each round, buyers were randomly rematched with sellers.[3]

For the experiment I chose the following parameters. The value of the product to the buyers was

$$v = \begin{cases} v_h = \$2.60 \text{ w.p. } 1/2 \\ v_l = \$1.20 \text{ w.p. } 1/2 \end{cases}$$

[3] I ran fewer treatments than I at first expected, and failed to balance the order of treatments. In four sessions, agents started with low verification cost before proceeding to high verification cost, and subjects played the treatments in the opposite order in only two sessions. This means that differences in treatment response could correspond slightly with differences in treatment order.

And the sellers could set price

$$p \in \{p_l = \$0.60,\ p_h = \$2.00\}$$

In the low-cost treatment, the cost of information was $c = \$0.10$, and in the high-cost treatment it was $c = \$0.30$. Subjects were paid for every choice in the experiment, but sellers were only paid half of the face value of their earnings so that average earnings were similar between buyers and sellers. Subjects were informed that this would be the case at the beginning of the experiment.

5 Results: Equilibrium Selection

There are two main possibilities for what could happen theoretically. On the one hand, buyers could refuse to buy any product at the high price. This would force sellers to always set the low price. Consumers would win out, but prices would convey no information about quality.

On the other hand, buyers may believe that expensive products are generally high-quality. If this were true, it could make it worthwhile to buy or at least check the quality of a high-price product. In this case, buyers' beliefs can become self-fulfilling. High-quality firms will know they can sell at a high price, while low-quality firms will not be so sure, and will mix between the two prices. Since high-quality firms always set the high price while low-quality firms mix, buyers turn out to be correct that prices convey information.

In the data, high-quality sellers are convinced that they can sell at a high price and they are correct. Buyers almost never reject an expensive product without at least verifying the quality first. As a result, high-quality sellers have a greater incentive to set the high price than low-quality sellers. The difference in how likely it is for high-quality and low-quality sellers to set the high price means that prices convey information to buyers.

Figure 4 shows how likely each type of seller is to set the high price in the various Nash

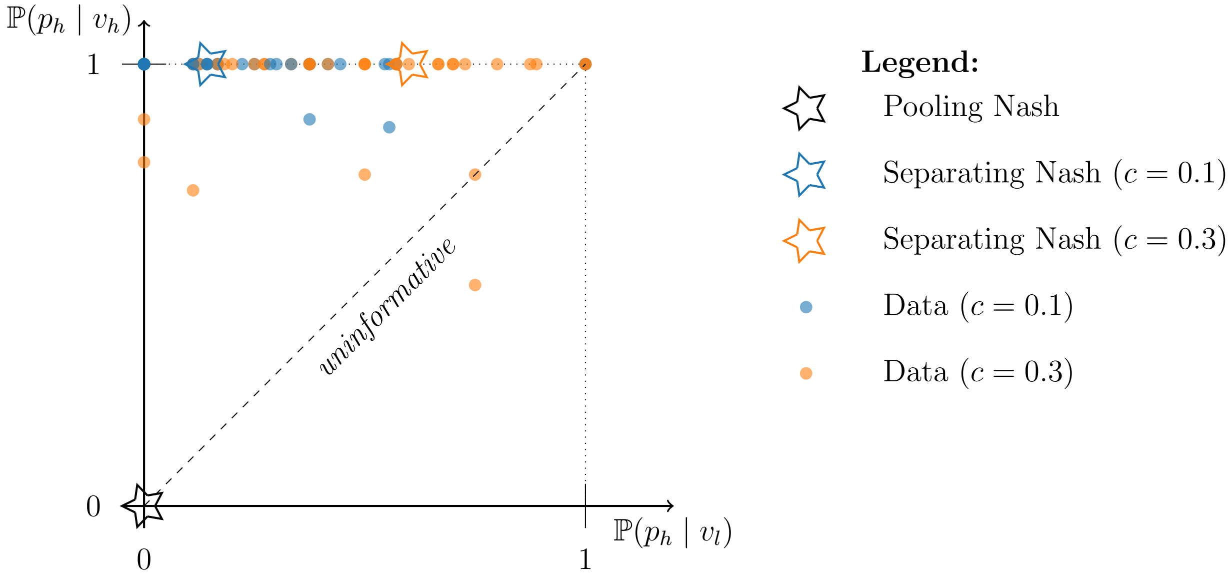

Figure 4: Pooling Nash is *not* Selected

Notes: Data is by-subject. Contrary to the pooling Nash equilibrium, almost all subjects are more likely to set the high price when they have a high-quality product, demonstrating that information is conveyed through the arbitrage channel in Grossman-Stiglitz (1980).

equilibria as well as in the data. The data overwhelmingly select the informative equilibrium discussed by Grossman and Stiglitz. Not only do buyers expect prices to convey information, but this expectation is self-fulfilling, and prices actually do convey information endogenously.

The diagonal line on figure 4 represents strategies where a firm will choose the same pricing decision whether it has high- or low-quality products. If both types choose the same prices, buyers cannot learn anything about quality from observing prices, so this diagonal also represents seller strategies that are *uninformative*–conveying no information to buyers about quality.

The data is firmly above this diagonal. High-quality products are consistently priced higher, on average, than low-quality products. Buyers understand this and generally buy or get information when facing a high price. If prices conveyed no information, buyers would simply walk away until the firm offers a low price. So prices provide meaningful information, inducing buyers to make choices they would not have made under their prior beliefs.

6 Results: Behavioral Deviation from Theory

Informational efficiency, (insofar as it is possible in the Grossman-Stiglitz paradox), is a lot to ask of buyers and sellers. The rational buyer must sometimes buy a product blindly, without checking the quality. They need to be quite precise themselves, and have perfect confidence in the precision of their fellow buyers, whose choices determine how trustworthy the market prices will be.

	$c = 0.1$		$c = 0.3$	
	Data	Mixed Nash	Data	Mixed Nash
$\mathbb{P}(p_h\|v_h)$	0.99	1.00	0.96	1.00
$\mathbb{P}(p_h\|v_l)$	0.23	0.14	0.46	0.60
$\mathbb{P}(B\|p_h)$	0.20	0.30	0.32	0.30
$\mathbb{P}(I\|p_h)$	0.78	0.70	0.58	0.70
$\mathbb{P}(B\|p_l)$	0.94	1.00	0.90	1.00
$\mathbb{P}(I\|p_l)$	0.04	0.00	0.05	0.00

Figure 5: Mixed Nash Generally Fits the Data

Notes: Generally, data aligns with the Grossman-Stiglitz Paradox. The biggest deviations from theory are: (1) Buyers get less information when it is more expensive, even though the Nash predicts no change in verification, and (2) low-quality sellers respond less strongly to the change in verification cost than they should according to the Nash theory. Both these facts are explored below, and predicted by quantal response.

It is as if a \$20 bill is left on the sidewalk. The well-trained economist just walks by, because if it were a *real* \$20 bill, someone else would already have taken it. Assuming the others walking by are perfectly mixing between sometimes checking and sometimes ignoring the potential \$20, the economist is justified in being indifferent to checking or just walking by. [4]

But this rigid calculus may not be intuitive for actual consumers. If their faith in others' rationality wavers, and the cost of bending down to check the \$20 bill is low, they may verify for themselves more often than is strictly necessary. Or, if checking the quality is quite difficult, they may place too much weight on their fellows and fail to check even when

[4] I heard this example from Maxim Engers; I've since learned its origin is older, but could not find a source for it.

they should. These behavioral tendencies could skew market prices away from the efficient markets prediction.

These behavioral tendencies can be seen from the choice probabilities in figure 5. When the cost of verification is low ($c = \$0.10$), buyers verify more than the Nash equilibrium predicts, not placing enough trust in the market. When the cost of verification is high ($c = \$0.30$), buyers verify less than the Nash equilibrium predicts, placing too much faith in the market. [5]

In the theory, buyers verify just as much when verification is expensive as when it is cheap. This is not intuitive for buyers in the lab. It happens in the theory because sellers respond steeply to the change in verification cost, even though this cost does not affect the seller payoff directly. When the verification cost increases, the low-quality seller tries to cheat the buyer more, so the buyer has a greater incentive to verify. This increased incentive to verify perfectly offsets the increased cost of verification, so that the buyer continues to verify at the same rate as before.[6]

Figure 6 shows the best responses of a low-quality seller and a buyer, for both high and low verification cost, along with session-level data.[7] While the perfectly vertical and horizontal best-response lines indicate that changes in information cost will lead to large differences in seller behavior but no change in buyer behavior, actual data shows that agents behave more smoothly. If the agent response functions were smoother, as would be the case if their were some noise in the agents' decisionmaking, we could rationalize (1) that buyers

[5]This should not be confused with buyers getting less, or more, information than is optimal given empirical seller choices. Since mixed Nash equilibria are unstable, if sellers deviate slightly from the equilibrium choice probabilities, buyers who optimally respond to those deviations will move *further* from Nash behavior, not closer to it. Here, I examine differences between actual behavior and equilibrium behavior, but it is important to remember that this is different from (in fact, opposite to) empirical optimality of subjects given opponents' out-of-equilibrium behavior when mixed Nash equilibria are involved.

[6]This is the standard intuition of a mixed Nash equilibrium. It is similar to how, in a game of rock-paper-scissors, if a third party offers an additional incentive to one player to play scissors (regardless of whether they win or lose), that player will not adjust their choice probabilities at all. Instead, their *opponent* will respond by playing rock more often. This increases the cost of playing scissors so that it is once again equal to the benefit of playing scissors, and the player remains indifferent.

[7]Again, since this is just a slice of the action space, other choice probabilities are fixed at their levels in the mixed Nash equilibrium.

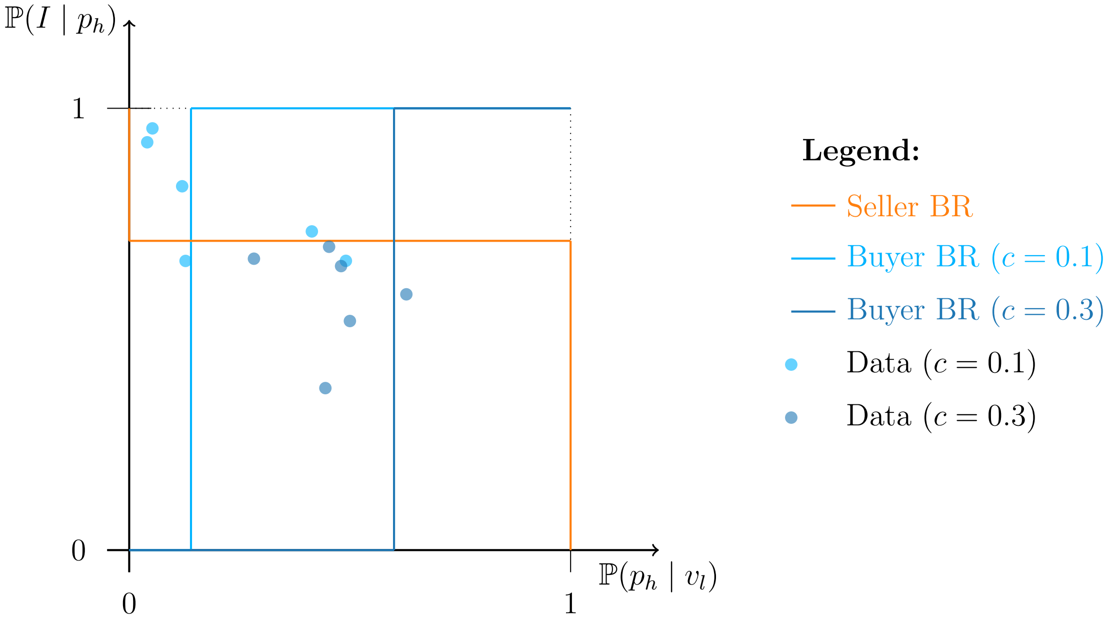

Figure 6: Subject Responses are Smoother than Theory

Notes: While data is fairly close to theoretical predictions, it is clear that the data for $c = 0.1$ (cyan) is higher than the data for $c = 0.3$ (blue), indicating buyers get less information when it is more expensive. Less obvious, but true, is that the data for the two treatments is closer together horizontally than the two Nash equilibria, indicating that sellers do not respond as strongly to the change in information cost as theory predicts.

respond to an increase in information cost by getting information less frequently, and (2) that sellers do not respond as drastically to the information cost as they would in the Nash equilibrium. Both of these facts are predicted by quantal response.

In a quantal response equilibrium, agents make small, mean-zero errors when evaluating which actions give them the highest expected payoffs. When two actions yield similar expected payoffs, agents play them with similar likelihood. When one action yields a much higher payoff than another, agents will play the better action much more frequently. Agents understand that these errors occur (in themselves and others) and respond accordingly, leading to an equilibrium where beliefs remain consistent, although their behavior sometimes deviates from their best responses.

Thus, if agents' expected payoffs are smooth functions of their opponents' choice probabilities, their own choice probabilities will also be smooth functions of their opponents behavior. Low-quality sellers will not suddenly decide to cheat buyers once buyers are verifying less than a certain threshold. And buyers will not suddenly decide to always verify when sellers begin cheating them more than a threshold amount.

When the cost of verification increases, buyers will respond gradually by buying less information. In turn, sellers will respond to buyers by cheating them slightly more often. Contrary to Nash, buyers will respond to changes in information cost, while sellers respond to those changes less drastically. This can be seen from quantal response curves plotted in figure 7. While quantal response does not always predict the levels correctly (in particular, sellers in the data are more hesitant to cheat buyers than quantal response predicts), it does predict the dynamics of how choice probabilities change as the cost of verification changes, and how those dynamics differ from Nash theory.

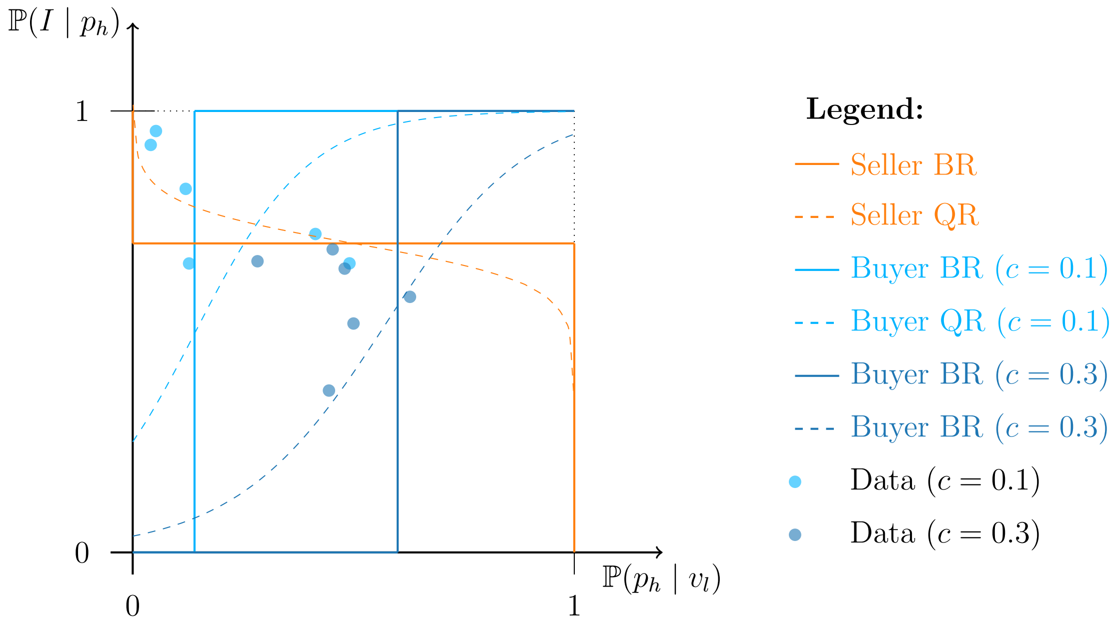

Figure 7: QRE Explains Differences in Verification and Cheating

Notes: This is the best-fit QRE (precision = 22). The QRE is fit just to this slice of the strategy space. The other decisions (seller's pricing decision when high-quality, and buyers' decision when facing a low price, etc.) are simple enough that agents are quite close to Nash. Quantal response does not predict the level of cheating well; sellers cheat less than predicted (perhaps due to an aversion to lying). However, quantal response does predict the differences across the two treatments: buyers respond more to changes in information cost and sellers respond less to those changes than Nash suggests.

The extent to which sellers are trying to cheat buyers determines the informativeness of prices. If low-quality sellers are constantly trying to price similarly to high-quality sellers, then prices will not convey much information about quality. Since behavioral noise smooths the relationship between seller choice probabilities and information cost, it ultimately smoothes the relationship between information cost and the informational content of prices. The Grossman-Stiglitz paradox asserts that prices will be more informative when arbitrage is cheap and less informative when arbitrage is expensive, so that the return to arbitrage always perfectly offsets its cost. That relationship still exists but is dampened by behavioral noise, as seen in figure 8.

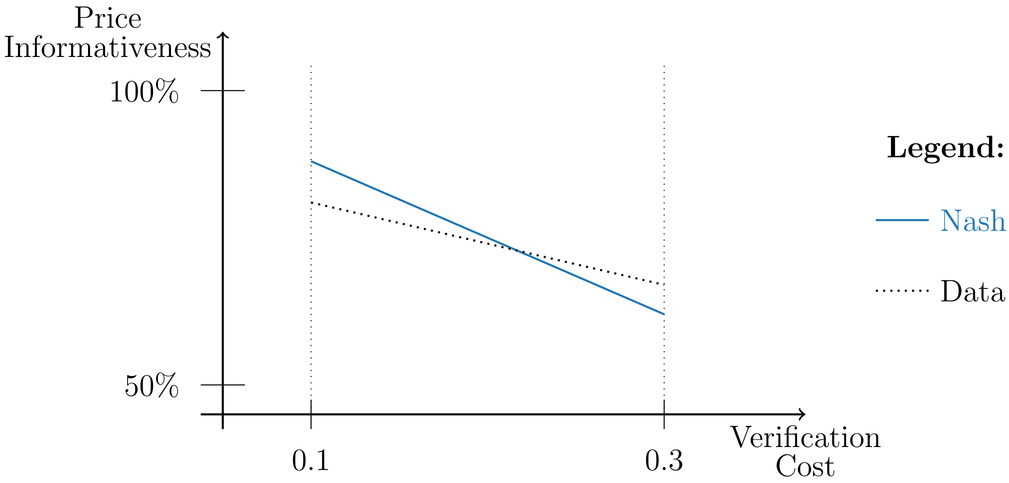

Figure 8: Cost and Benefit of Verification

Notes: Informativeness is measured as the unconditional probability that the quality can be guessed from the price. So 50% is completely uninformative: no better than random guessing. 100% means the quality is known for certain after the price is observed. These lines are just interpolated from the identified endpoints.

While quantal response can provide some intuition for the dynamics, we can test behavioral tendencies using a permutation test. Under the null hypothesis, buyers follow the Nash and do not adjust their behavior as the cost of verification changes. Since the verification cost should not matter (under the null hypothesis), permuting the labels of the data–which data points have come from which treatment–should not matter. If, instead, the observed data involves treatment differences more severe than permuted versions of the data, we can conclude that verification cost does affect buyers' choice probabilities. In the same way, we can test whether low-quality sellers do indeed change their behavior less than they should in the theory, and whether informativeness therefore adjusts less. The results are given in figure 9.

Test	p-value
Verification decreases with cost	$2/64 = 0.031$
Cheating increases less than Nash	$3/64 = 0.047$
Informativeness decreases with cost	$1/64 = 0.016$
Informativeness decreases less than Nash	$6/64 = 0.094$

Figure 9: Non-parametric Permutation Tests

Notes: With six paired observations, there are 64 possible permutations of the data. The p-value is the fraction of these permutations that produce a test statistic at least as extreme as the one observed in the actual (unpermuted) data. In each case, the test statistic is the average difference between the two treatments.

We can conclude that behavioral noise leads buyers to respond too much to the cost of verification, and leads sellers to respond too little. Thus, while prices are more informative when information is cheap, the effect appears to be smaller than Nash. When information is cheap, agents purchase more information than the Nash predicts, and when information is expensive, agents buy less information than Nash predicts.

7 Conclusion

Overall, in the lab, markets are remarkably efficient in terms of how much information prices convey, even with a relatively small number of untrained subjects. Agents' choice probabilities are very close to those predicted by Grossman and Stiglitz. This demonstrates that agents do not simply *believe* that prices convey information in real-world markets and take those beliefs into the lab, but that prices actually *do* convey information endogenously in the lab. Specifically, prices can convey information via the Grossman-Stiglitz channel of costly verification.

Nevertheless, there are still some behavioral deviations from theory, when subjects trust the market too much or too little. These are intuitive because they would arise if subjects had smooth best-response functions rather than the infinitely steep best-response functions of a rational agent. When verification is easy, subjects verify more than Nash, failing to realize that everyone else can also verify more easily and the market adjusts. When verification is hard, subjects verify less often than Nash, failing to realize that verification is also difficult for everyone else, and they should trust the market prices less.

Appendices

A Full extensive form

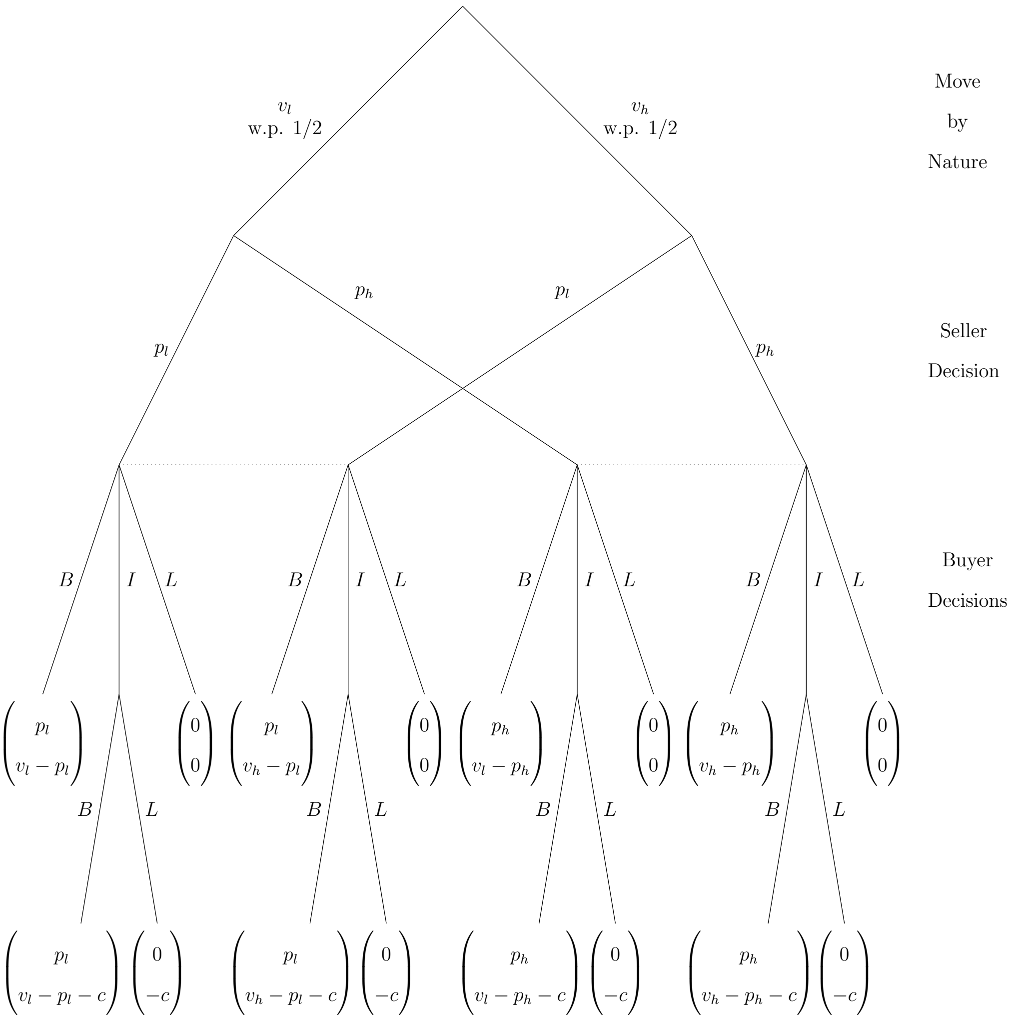

Figure 10: Extensive Form of the Game

B Derivation of equilibria

B.1 Pure strategy equilibrium

There is no perfectly separating equilibrium in this game. If the seller were to choose the high price when quality is high and the low price when quality is low, consistent beliefs imply that the buyer would know in equilibrium that a high price implies high quality. Thus, the buyer would no longer need to pay the information cost because they can learn quality costlessly from the price. But if the buyer is not paying the information cost, there is an incentive for the seller to charge the high price even when the quality is low, since consumers will expect quality to be high and will be fooled into buying the product. Thus, there is no separating equilibrium where the seller prices according to their quality level. Consistent beliefs also rule out reverse type dependence; it is not an equilibrium for the high quality seller to set the low price and the low quality seller to set a high price.

It could be that the seller will set the same price whether their quality is high or low. In this case, consumers will buy only if the expected quality is greater than the price. I have chosen the possible prices so that

$$p_l < v_l < \frac{1}{2}(v_h + v_l) < p_h \tag{1}$$

and this ensures that buyers will not buy if the seller is always setting the high price, but will buy if the seller sets the low price. So there is no equilibrium where the seller always sets the high price, but there is an equilibrium where the seller always sets the low price.

For this to be an equilibrium, the seller must not prefer to deviate to the high price, and thus we need the buyer to choose not to get information and not to buy if they were to observe the high price. This can be ensured by buyer beliefs. Let μ_h denote the probability that the seller is high quality given that they set the high price. Since the high price is off the equilibrium path, μ_h is unconstrained. That is, if participants reach an equilibrium where the high price is never set, buyers can reasonably assume anything about the expected quality a firm would have if they were to set the high price. If the payoff of not getting information and not buying is higher than both getting information and buying without information, we have:

$$\mathbb{E}[u(L)|p_h] \geq \mathbb{E}[u(B)|p_h] \qquad \mathbb{E}[u(L)|p_h] \geq \mathbb{E}[u(I)|p_h] \tag{2}$$

$$0 \geq \mu_h(v_h - p_h) + (1 - \mu_h)(v_l - p_h) \qquad 0 \geq \mu_h(v_h - p_h) - c \tag{3}$$

For this to hold, the buyer must believe that, conditional on setting the high price, the seller is unlikely enough to be high quality that expected quality does not exceed the price, and that the expected benefit of information does not exceed the information cost.

$$\mu_h \leq \min\left\{\frac{p_h - v_l}{v_h - v_l}, \frac{c}{v_h - p_h}\right\} \tag{4}$$

As long as the buyer believes a deviation to the high price is sufficiently likely to occur when the seller is low-quality, pooling at the low price can be sustained as an equilibrium.

B.2 Mixed Equilibrium

In this equilibrium, the seller always chooses the high price when they have high quality. When the seller has low quality, they mix between the two prices, choosing the high price with probability $\mathbb{P}(p_h|v_l)$. Thus, if the buyer sees the high price, they believe the seller is high-quality with probability $\mu_h = 1/(1 + \mathbb{P}(p_h|v_l))$. In order for the seller to mix when they have low quality, it must be that the expected profit from setting the low price is equal to the expected profit from setting the high price. This implies that when the buyer observes the high price, they only buy without getting information a fraction of the time:

$$\pi_h(p_h) = \pi_l(p_h) \implies p_l = \mathbb{P}(B|p_h)p_h \implies \mathbb{P}(B|p_h) = \frac{p_l}{p_h} \tag{5}$$

and thus the buyer chooses to get information with corresponding probability $\mathbb{P}(I|p_h) = \frac{p_h - p_l}{p_h}$. Because the buyer's choice probabilities must make the seller indifferent in the mixed Nash, the buyer's probability of getting information is independent of the cost of information. Similarly, the seller's choice probabilities must make the buyer indifferent. For the buyer to want to mix between buying without information and getting information when observing the high price, it must be that their expected payoffs are the same:

$$\mathbb{E}[u(B)|p_h] = \mathbb{E}[u(I)|p_h] \implies \mu_h v_h + (1 - \mu_h) v_l - p_h = \mu_h (v_h - p_h) - c \tag{6}$$

and this is true when the seller's choice probability is

$$\mathbb{P}(p_h|v_l) = \frac{c}{p_h - v_l - c} \tag{7}$$

If the seller has low quality, they must set the high price more often when the effort cost rises so that the buyer will have a stronger incentive to get information and remain indifferent between getting information and buying without information. These comparative statics are typical of mixed equilibria: if information becomes more expensive, the buyer (who pays the higher information cost) still buys information at exactly the same frequency, and the seller (who does not pay the cost) changes their behavior.

In order for this equilibrium to exist, it must be that the payoffs of buying without information or getting information are no lower than the payoff of not getting information and not buying. This is true when the effort cost is sufficiently low:

$$\frac{c}{p_h - v_l - c} \leq \frac{v_h - p_h}{p_h - v_l} \implies c \leq \frac{(v_h - p_h)(p_h - v_l)}{v_h - v_l} \tag{8}$$

Chapter 2

Competition and Price Informativeness: An Experiment

Abstract

Consumers often rely on price as a guide to the quality of a product. If consumers are unable to observe quality directly, price might not perfectly communicate product quality. In this context, competition might increase or decrease the informativeness of prices. I study how competition impacts the informativeness of prices theoretically and in a lab experiment. While theory leaves the question open, I find in the experiment that competition leads both high- and low-quality firms to decrease prices, but the price reduction is larger for high-quality firms who are more likely to price high in the absence of competition. Thus, prices become a less reliable guide to quality when there are more sellers in the market.

1 Introduction

Prices often contain information about the quality of products. Consumers are primed to expect prices to carry information, and will use the price as a guide to quality when quality is otherwise unknown (Leavitt 1954, McConnell 1968, review in Olson 1977). Nevertheless, prices do not do this job perfectly. Consumers continue to occasionally experience ex-post regret after buying a product of lower quality than expected. Or consumers may expend time and effort to learn about product quality prior to purchase, implying that the price is not, on its own, a fully reliable guide to quality.

Prices are fully informative of quality if consumers can predict the quality of a product from its price, as would be possible if high-quality products always sold for high prices and low-quality products always sold for low prices. But one main reason prices are not fully informative is that, if consumers are willing to buy a high-quality product at a high price, firms with low-quality products may be incentivized to 'pretend' to be high-quality by setting a high price. If the consumer cannot observe quality directly and uses the price as a guide, this mimicking strategy could be successful.

In this context, competition between firms becomes relevant. Competition generally lowers prices. But the effect of competition on price *informativeness* is theoretically unclear. It could be that competition drives down the prices of the low-quality firms, and thus makes price a more reliable guide to quality. But it might also be that competition drives down the prices of the high-quality firms and thus price does a worse job distinguishing between firms.

In this paper, I examine how competition between sellers affects how well prices function as signals of product quality. I use a simplified version of Janssen and Roy (2010), and consider the perfect Bayesian equilibria. I show how different patterns of equilibrium selection could lead to more or less informative pricing, and then implement the game in a lab experiment.

In the experiment, I find that competition generally decreases price informativeness. Competition drives down the prices set by both high- and low-quality firms, but the effect is largest on high-quality firms, who are most likely to set high prices in the absence of competition. Thus, the variance of prices decreases, and consumers are less able to deduce the quality of a product from its price.

In a standard signalling model, there are two types, high and low, and the high-type engages in some costly signalling behavior to prove their high type. Maybe this is a worker getting an education to prove their high ability, or a firm getting a costly certification of their product. The signalling behavior needs to be less costly for a high-type than for a low-type; that way the high-type can credibly demonstrate that they are a high-type by doing something the low-type would never want to do. The recipient of the signal believes the message because they know that the signal would be so costly for a low-type that the low-type would not want to mimic the signal, even if it meant they could masquerade as a high-type.

In my simplified version of Janssen and Roy, one or more firms have either high or low quality, and choose to set either a predetermined high price or a predetermined low price. Buyers see the price, but not the quality of the firms. When signalling occurs through the sellers' price choices, it is harder to sustain separation because the signalling behavior is more suspect. If a firm tries to signal their high quality by setting a high price, it is more difficult to convince buyers, since buyers know that *any* firm, high- or low-quality, would prefer to sell at a high price. This is the additional wrinkle in a price signalling game compared to a generic signalling game.

In order for signalling to be credible, it must be costly, and this means a firm must be less likely to sell when setting a high price than setting a low price. One way this could happen is if low-quality firms occasionally set the high price; then buyers would be suspicious of the high price and sometimes not buy high-priced products.

If this is happening and then an increase in competition drives prices down, the relevant question is: whose prices are driven down? The reason prices are not fully informative is because low-quality firms sometimes set high prices, so if competition drives down low-quality firms' prices, prices become more informative. But if competition drives down the prices of high-quality firms, it becomes harder to distinguish between firms based on prices, so prices become less informative.

This work contributes to the literature in two ways. The model can be fully solved out for the symmetric equilibria, and to my knowledge, the solution is new. This model provides a simple environment in which to understand price signalling. This work contributes to the experimental literature by examining how prices might convey information endogenously and how that role is affected by competition. There is a large experimental literature that examines how well prices function as guides to product quality, but none examine prices conveying information endogenously. These are generally older papers where the rational expectations of buyers are not considered. As yet, I know of no experiments that examine endogenous price signalling.

2 The Model

The model consists of n ex-ante-identical firms. Each firm is independently equally likely to be high-quality or low-quality. After realizing its quality, each firm simultaneously chooses a price $p \in \{p_L, p_H\}$. There is one buyer, who observes the price set by each firm but cannot observe which firms are high-quality and which are low-quality. After seeing the vector of prices, the buyer decides whether to buy one of the high-priced products, one of the low-priced products, or none. If the buyer chooses not to buy, they receive a payoff of zero, as do all the firms. If the buyer chooses to buy one of the products, they receive a payoff of $v - p$ where p is the price of the product they chose to buy, and v is the value of the product. $v = v_H$ if the firm is high-quality and $v = v_L < v_H$ if the firm is low-quality. The firm whose product sold receives $p - c$ where p is the price the firm set and c is the cost of production. If the firm is high-quality, $c = c_H$, but if the firm is low-quality, $c = c_L < c_H$. All other firms receive a payoff of zero. The timing of the game is given in figure 1.

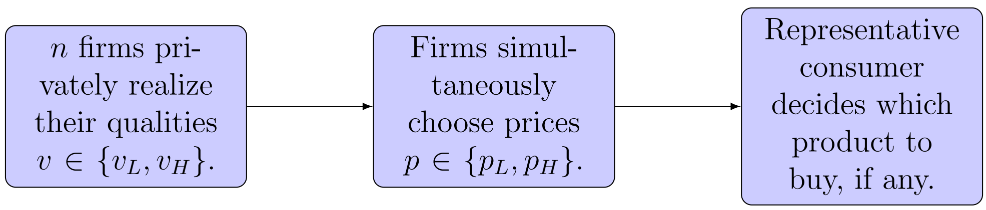

Figure 1: Timing of the Game

The simplification of sellers' pricing decision to a binary choice between an exogenously set p_L and p_H vastly simplifies the model; this is my main simplification from the larger model in Janssen and Roy (2010). This simplification is necessary to solve out for the symmetric perfect Bayesian Nash equilibria, but also to make it easy for subjects to understand the experiment, and to cleanly analyze the resulting choice data. Nevertheless, this is a huge simplification, and I attempt to examine the robustness of the empirical results by also running treatments where subjects choose from a price grid.[1]

[1] Without being able to fully solve the model with a continuous price choice, it's difficult to see what may be left out by this simplification. In the broader model where firms can choose any price, single-crossing

A firm's pure strategy involves a price to set if the firm is high-quality and a price to set if the firm is low-quality. Each potential mixed strategy can be denoted by

$$s = \begin{cases} (p_H, p_H); \text{ w.p. } \mathbb{P}_1 \\ (p_H, p_L); \text{ w.p. } \mathbb{P}_2 \\ (p_L, p_H); \text{ w.p. } \mathbb{P}_3 \\ (p_L, p_L); \text{ w.p. } \mathbb{P}_4 \end{cases}$$

where (a, b) denotes the pure strategy of setting $p = a$ when high-quality and $p = b$ when low-quality. A pure strategy for the buyer is a choice to buy either a high-priced product, a low-priced product, or nothing, for each possible vector of prices that could be observed. I further limit the analysis to fully symmetric equilibria, and the solution concept is the perfect Bayesian Nash equilibrium.[2]

I consider parameter values in which a low-quality product is worth buying at the low price and the high-quality product is worth buying at the high price, but a low-quality product is not worth buying at the high price:

$$p_L < v_L < p_H < v_H$$

still holds, which means that there exists a $\hat{p}$ such that all prices set in equilibrium by high-quality firms are weakly higher than $\hat{p}$, and all prices set in equilibrium by low-quality firms are weakly lower than $\hat{p}$. In a perfect Bayesian equilibrium, the buyer will know the quality exactly for any price besides $\hat{p}$ (as long as it is on the equilibrium path). Perhaps the most significant difference between the larger model and my simplified version is that it is possible to achieve full separation in the larger model. Even with only a single firm, full separation is possible if the low-quality firm chooses v_L and the high-quality firm chooses v_H. This equilibrium involves the buyer playing a weakly dominated strategy, which I find very unlikely in the lab, but with more firms, separation remains possible in more complicated ways.

[2]This means firms of the same type play the same strategy. It also means, since the buyer only observes prices and has no other information about the firms, the buyer cannot distinguish between firms setting the same price. If the buyer chooses to buy a low-priced product (for instance) and there are multiple firms setting the low price, the buyer chooses uniformly randomly among them, so that each sell an equal fraction in expectation.

I preclude firms from pooling at the high price by assuming that

$$p_H > \mathbb{E}[v]$$

This means low-quality firms cannot *always* cheat the buyers without buyers eventually deciding to stop buying the expensive products. I assume that consumers prefer high-quality products at the high price to low-quality products at the low price.

$$v_H - p_H > v_L - p_L$$

If this were not true, at least with sufficient probability, then buyers would always prefer cheap products to expensive products, regardless of quality, and consumers would never regret their purchase. Lastly, I allow both firms to have positive profit margin when setting the low price

$$c_H < p_L$$

This is because, in this model, the effect of competition on informativeness comes down to whether high-quality or low-quality firms decrease their price the most in response to competition. If this assumption were not true, high-quality firms would never consider a price other than p_H because it would lead to losses.

2.1 Equilibria

Consider the buyer's problem after observing a vector of prices $\vec{p}$. Given the observed vector of prices, the buyer has beliefs about the likelihood that each firm is high-quality versus low-quality. I denote the buyer's belief that firm j is high-quality after observing $\vec{p}$ by

$$\mu_j(\vec{p}) \equiv \mathbb{P}(v_j = v_H | \vec{p})$$

Since consumers cannot distinguish between firms setting the same price, and only choose whether to buy a high-priced or low-priced product (or none), the relevant beliefs are

$$\mu_H(\vec{p}) \equiv \mathbb{P}(v = v_H | p = p_H, \vec{p})$$
$$\mu_L(\vec{p}) \equiv \mathbb{P}(v = v_H | p = p_L, \vec{p})$$

Further narrowing the problem by focusing on symmetric equilibria hugely simplifies beliefs by reducing the dimensionality from $2n$ down to 2, everywhere along the equilibrium path.

Proposition 1. *In any symmetric equilibrium where* $\mathbb{P}_1 \neq 1$*, and* $\mathbb{P}_4 \neq 1$*,*

$$\mu_H(\vec{p}) \equiv \mu_H, \;\; \mu_L(\vec{p}) \equiv \mu_L$$

This is obvious since, for any firm setting the high price, $\mu_H = (\mathbb{P}_1 + \mathbb{P}_2)/(2\mathbb{P}_1 + \mathbb{P}_2 + \mathbb{P}_3)$ and for any firm setting the low price, $\mu_L = (\mathbb{P}_3 + \mathbb{P}_4)/(\mathbb{P}_2 + \mathbb{P}_3 + 2\mathbb{P}_4)$ by Bayes' Rule. But the requirement that the equilibrium be symmetric is necessary. Otherwise, there could be, for instance, one firm that always sets the high price regardless of quality and another firm that is type dependent, setting the high price when high-quality and the low price when low-quality. In this case, if the buyer observes $\vec{p} = (p_H, p_L)$, they know the high price was set by the firm that always sets the high price, and thus $\mu_H(\vec{p}) = 1/2$. But if the buyer observes $\vec{p} = (p_H, p_H)$, then one of the high prices comes from a surely high-quality firm and one comes from a firm with $1/2$ chance of high-quality, so $\mu_H(\vec{p}) = 3/4$. So if firms are not symmetric, the whole price vector can matter for buyer beliefs, and the buyer can have different beliefs for different amounts of high and low prices in the price vector.

The caveat that $\mathbb{P}_1 \neq 1$ and $\mathbb{P}_4 \neq 1$ is also necessary. If one of the two prices is never set, then most of the price vectors will never be reached. If the situation where the buyer observes $\vec{p}$ is off the equilibrium path, then beliefs are not constrained by observed behavior, so the buyer can have any beliefs $\mu_H(\vec{p})$ and $\mu_L(\vec{p})$, and thus $\mu_H(\vec{p})$ will not necessarily be the same as $\mu_H(\vec{p'})$.

Continuing with the assumptions of proposition 1, consider the buyer's expected utility from each of their possible strategies. If the buyer buys a high-priced product (strategy B_H), they receive an expected payoff of

$$\mathbb{E}u_B(B_H) = \mu_H v_H + (1 - \mu_H)v_L - p_H$$

If they buy a low-priced product (strategy B_L), they receive

$$\mathbb{E}u_B(B_L) = \mu_L v_H + (1 - \mu_L)v_L - p_L$$

and they receive 0 if they buy nothing.

Notice that, because $p_L < v_L < v_H$, the buyer always prefers buying a low-priced product to buying nothing. The buyer will prefer buying a high-priced product to buying nothing if high-priced products are sufficiently likely to be high-quality:

$$\mathbb{E}u_B(B_H) \geq 0 \implies \mu_H \geq \frac{p_H - v_L}{v_H - v_L}$$

and the buyer will prefer a high-priced product to a low-priced product if the high-priced product is sufficiently *more* likely to be high-quality than the low-priced product:

$$\mathbb{E}u_B(B_H) \geq \mathbb{E}u_B(B_L) \implies \mu_H \geq \frac{p_H - p_L}{v_H - v_L} + \mu_L$$

We can partition the space of buyer beliefs into regions based on what strategy the buyer will play, in figure 2.

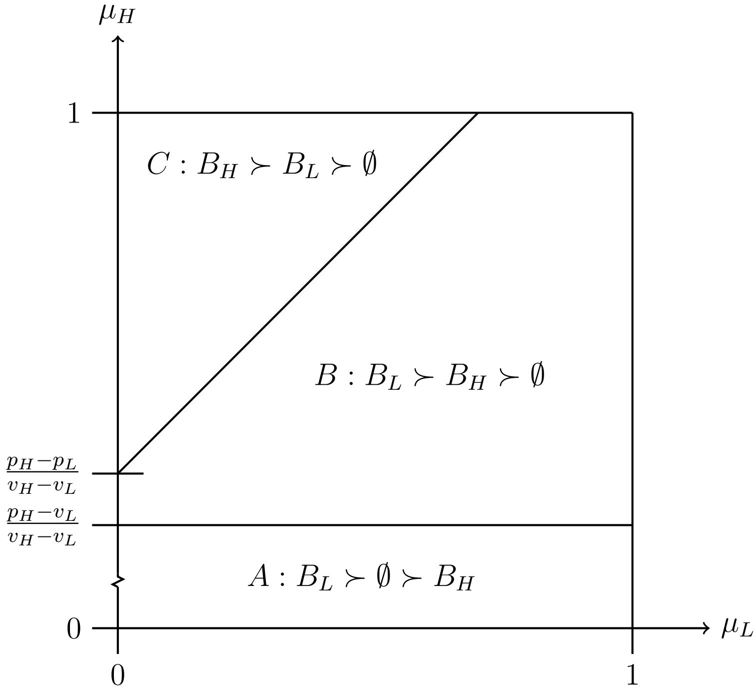

Figure 2: Buyer Belief Space

Notes: This figure represents the space of buyer beliefs: μ_H is the buyer's belief that a high-priced product will be high-quality, and μ_L is the buyer's belief that a low-priced product will be high-quality. Moving from the bottom right to the top left of the figure increases the perceived value of high-priced products relative to low-priced products.

Moving from the bottom right to the top left of figure 2, high-priced products become increasingly more attractive to the buyer. In region A, the buyer will buy a low-priced product if one exists, but if not, the buyer will buy nothing. In region B, the buyer prefers to buy a low-priced product, but if none exists, the buyer will buy a high-priced product. In region C, the buyer prefers a high-priced product and only buys a low-priced product if no high-priced product exists.

In region A, the buyer does not believe the signal, and will only buy a low-priced product. If this is true, it is in each firm's best interests to set the low price, regardless of the quality of its product. Thus, pooling at the low price can be sustained as an equilibrium as long as μ_H (which is free, since high prices do not occur in equilibrium) is sufficiently low. In this equilibrium, prices convey no information about the quality of the products.

In region C, the buyer believes the signal and strictly prefers to buy a high-priced product. But this means that a firm is always more likely to sell by setting a high price than a low price. Signalling has no cost—it benefits sellers by making them more likely to sell *and* yields higher profits when they do sell. All firms will want to set the high price, regardless of quality, and the signalling cannot be credible in equilibrium.[3]

This feature is common to all equilibria, both in this model and in many similar models where quality is unknown to buyers: In any equilibrium, the buyer must be sufficiently unlikely to buy the high-priced product. In order to sustain price signalling, it must be that setting the high price has a cost that counterbalances the obvious benefit of the increased profit margin. If the buyer buys a high-priced product with large enough probability, then low-quality firms will be better off setting the high price than the low price. This "cheating" from the low-quality firms means that buyers were wrong to buy the high-priced product.

In region B and its boundary, the buyer may sometimes buy a high-priced product, but is potentially less likely to buy a high-priced product than a low-priced product. Thus, signalling quality by setting a high price is costly. But in order for signalling to convey information in equilibrium, it must be that the signalling behavior is specifically *more* costly for low-quality firms than high-quality firms.

This can happen because of the differences in the cost of production between high- and low-quality firms. A low-quality firm makes $p_L - c_L$ from selling a product at the low price, and $p_H - c_L$ from selling a product at the high price. Thus, for a low-quality firm to choose to set the low price, they must be at least

$$\frac{p_H - c_L}{p_L - c_L}$$

[3]To see this, suppose a seller's competitors are setting the high price with probability x, and the low price with probability $1-x$. If the seller sets the low price, they will sell only if all other sellers set the low price, in which case they will share the expected surplus with the other $n-1$ firms and receive $(p_L - c)/n$. If instead the seller sets the high price, the worst that can happen is if all the other firms set the high price and the seller has to share the expected surplus, and receive $(p_H - c)/n$. Since the worst-case scenario when setting the high price is strictly better than the best-case scenario when setting the low price, all sellers should set the high price. But if all sellers are setting the high price, regardless of quality, then a high-priced product is just as likely to be low-quality as high-quality. This means its expected value, $(v_H + v_L)/2$, is lower than the price, p_H, and the buyer should not buy.

times more likely to sell at the low price than the high price. A high-quality firm has higher cost of production, and thus lower profit margins for a given price. For a high-quality firm to set the low price, they would need to be at least

$$\frac{p_H - c_H}{p_L - c_H} > \frac{p_H - c_L}{p_L - c_L}$$

times more likely to sell at the low price than the high price. This intuition leads to a single-crossing result.

Proposition 2 (Single Crossing). *Given a strategy of the buyer consistent with some beliefs μ_H and μ_L, and given a symmetric strategy for $n-1$ other firms, an individual firm's expected profit, conditional on a cost of production c, satisfies either*

$$\forall c \in [0, p_L],\ \mathbb{E}[\pi(p_H|c) - \pi(p_L|c)] > 0$$

or

$$\frac{\partial}{\partial c}\mathbb{E}[\pi(p_H|c) - \pi(p_L|c)] > 0$$

Stated another way, in situations where high- and low-quality firms differ in the prices they set, a high-quality firm always has a greater incentive to set the high price than a low-quality firm. Thus, if high-quality firms are mixing between the two prices (and therefore indifferent between them) low-quality firms will prefer to set p_L, and similarly, if low-quality firms are indifferent between the two prices, high-quality firms will prefer to set p_H. Figure 3 shows how this reduces the space of potential equilibrium seller strategies.

But we can immediately rule out pooling at p_H and pure type dependence, based on the reasoning from region C above. If both types set p_H, the expected value of a high-priced product is lower than it's price, and the buyer will not buy. Then firms will regret their strategy. If firms fully separate, then buyers will know in equilibrium that a high-priced product is always high-quality and a low-priced product is always low-quality. Since buyers

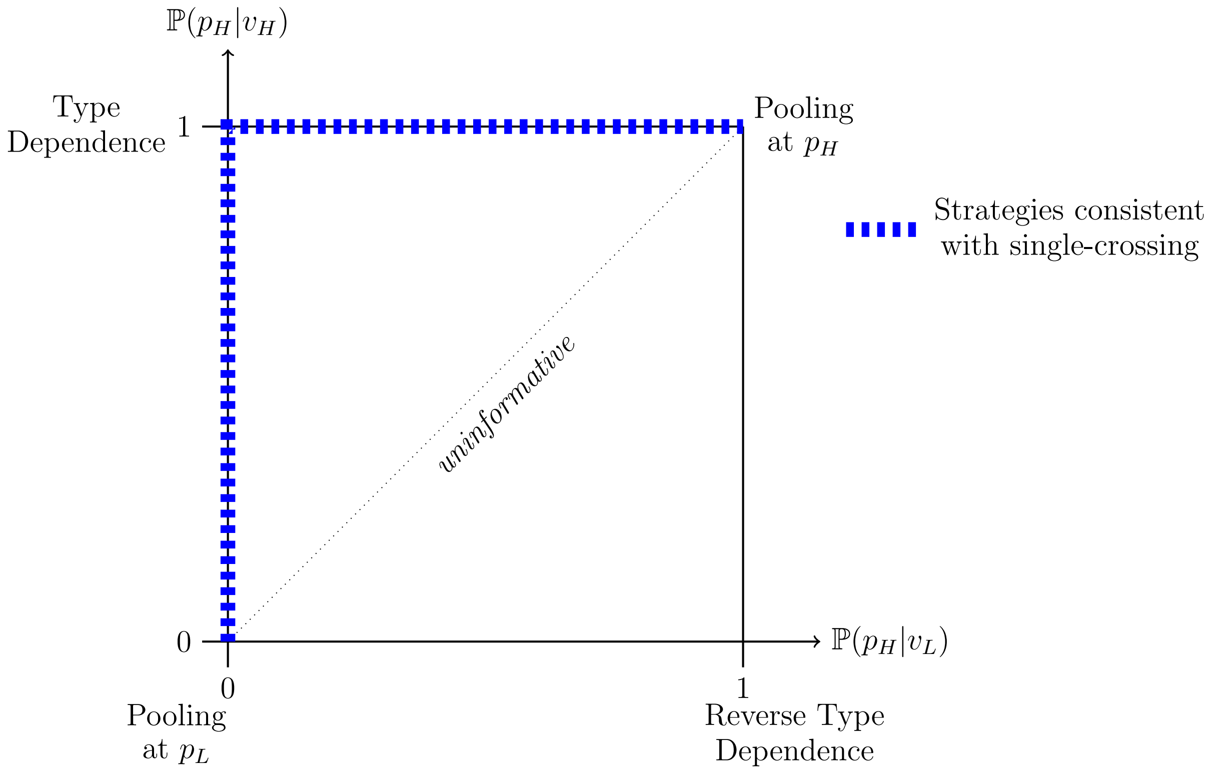

Figure 3: Seller Strategy Space

Notes: Each point (x, y) represents a seller strategy of setting p_h with probability x when low-quality and setting p_h with probability y when high-quality. Single-crossing rules out all but the highlighted strategies. Pooling at p_h and perfect type dependence are ruled out by consistency of buyer beliefs. Ultimately, the only possible equilibrium strategies are pooling at p_l (uninformative) and one or two mixed equilibria (partially informative).

prefer high quality at a high price to low quality at a low price, they will always opt for the high-priced product. But then the low-quality firms regret setting the low price.

So the potential equilibria involve pooling at the low price, or partially separating, either because the the high-quality firm occasionally sets the low price or because the low-quality firm occasionally sets the high price. All three of these types of equilibria turn out to be possible. In the pooling equilibria, prices convey no information, but prices are somewhat informative in the other equilibria. The next section looks at how the equilibria evolve as the number of firms increases.

3 Competition, Price Level, and Informativeness

It might seem intuitive that an increase in the number of firms would always weakly drive down prices, but this is not necessarily true. It is usually true that, given a strategy for the buyer, an increase in n increases the basin of attraction of p_L and shrinks the basin of attraction for p_H, but this is a result about sellers' best-response functions and not about equilibrium behavior. It can be that in a mixed equilibrium with increasing reaction functions among the firms, an increase in n makes firms set p_H more frequently in order to keep each other indifferent in equilibrium.

Nevertheless, a firm is most incentivized to set p_H when $n = 1$ and there are no competitors who could set p_L and undercut the firm. In this case, there is a unique informative equilibrium where a firm will certainly set the high price when high-quality, and will sometimes set the high price when low-quality as well. This leads to the following weaker statement about the relationship between competition and average price.

Proposition 3. *When $n = 1$, there is a unique informative equilibrium, and the probability that a firm will set the high price is at least as high in this equilibrium as in any equilibrium for any $n \in \mathbb{N}$.*

Figure 4 shows the price level in various equilibria as n increases. If agents get to the informative equilibrium when $n = 1$, increasing competition cannot raise expected prices relative to that benchmark. But a decrease in prices could mean an increase or a decrease in price informativeness, depending on whether the decrease in prices comes from high-quality or low-quality firms.

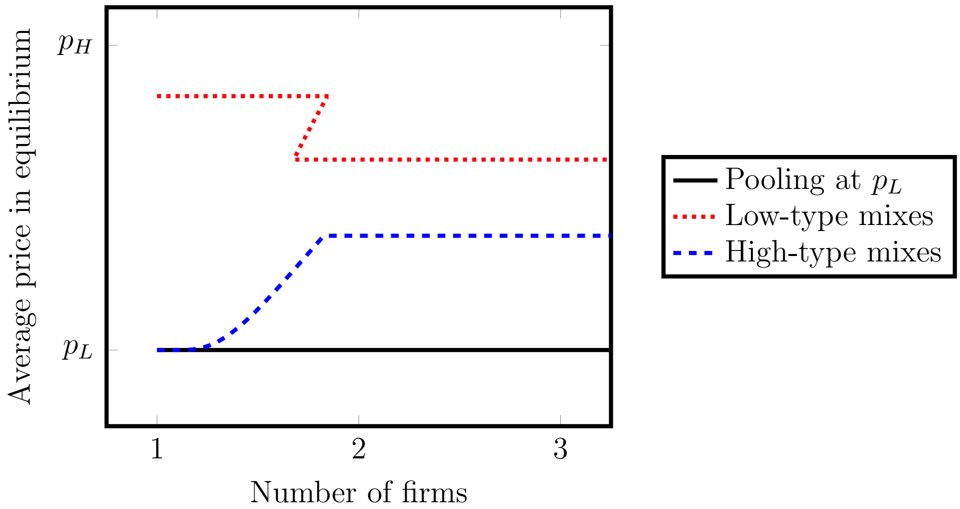

Figure 4: Average price as competition increases

Notes: This figure is based on the actual parameter values used in the experiment. Because firms are symmetric, the average price is simply $xp_H + (1-x)p_L$ where x is the unconditional probability of an individual firm setting the high price. For clarity, this figure graphs the equilibria as n varies continuously, even though, of course, only integer values of n make sense.

Suppose competition incentivizes low-quality firms to set the low price more often to stay competitive, but that the pricing strategy of high-quality firms is relatively unchanged. If this happens, buyer beliefs will update as a result. Since low-quality firms are setting the low price more frequently, when the buyer *does* see the high price, they infer that it is more likely to entail high quality. This makes a high-priced product more attractive to the buyer and means that the high-quality firms can continue to set the high price.

Alternatively, it might be that an increase in competition causes high-quality firms to set the low price more often as well, and this could decrease the informativeness of prices. Figure 5 shows different ways that increasing n could change equilibrium seller strategies and thus the informativeness of prices.

As $n \to \infty$, the buyer has access to at least one low-priced product with probability

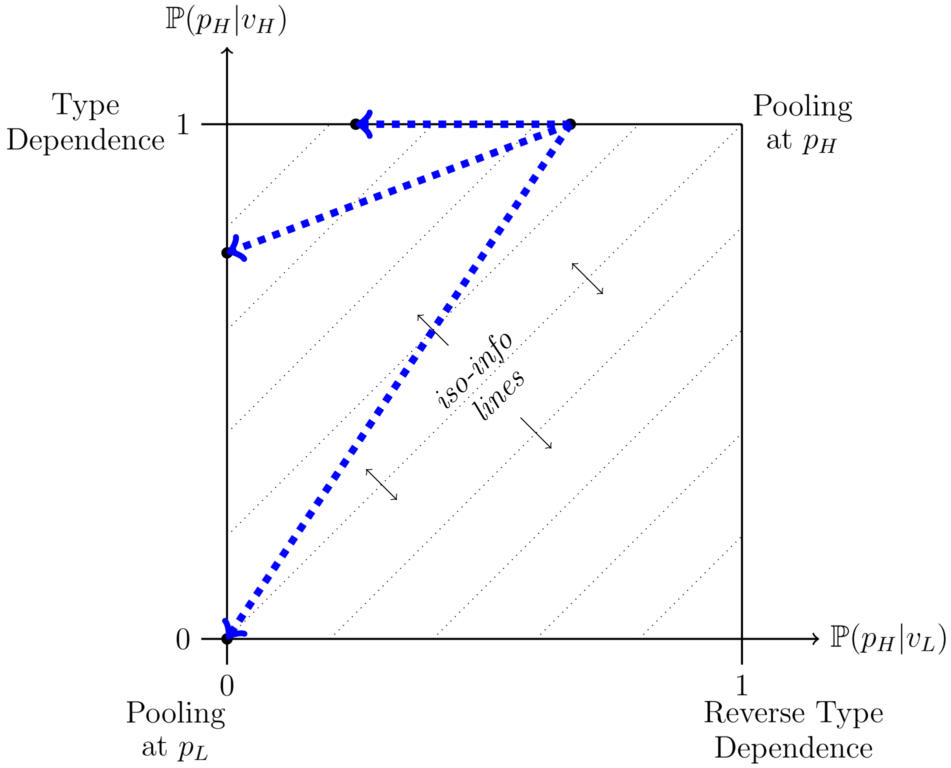

Figure 5: Potential Seller Strategies as n Increases

Notes: If subjects reach the informative equilibrium when $n = 1$, moving to $n = 2$ could either increase or decrease price informativeness, depending on equilibrium selection. This figure is based on actual parameter values used in the experiment.

approaching 1, since all equilibria involve each firm setting the low price with non-vanishing probability. The buyer cannot strictly prefer to buy a high-priced product, or all firms would set the high price (and that cannot be an equilibrium, as shown above). If the buyer always buys a low-priced product when it exists, then as $n \to \infty$, high-priced products will never be sold and firms will have to pool at the low price. The only alternative is for firms to make the buyer just indifferent between the two prices; any more information than that would lead the buyer to prefer the high price and could not be an equilibrium. Proposition 4 formalizes this, and figure 6 shows the informativeness of different equilibria as n increases.

Proposition 4. *As $n \to \infty$, equilibrium informativeness converges to either 1/2 (completely uninformative pricing) or*

$$\frac{3}{2} - \frac{v_H - v_L}{2(p_H - p_L)}$$

which is the maximum informativeness that can be sustained in equilibrium.

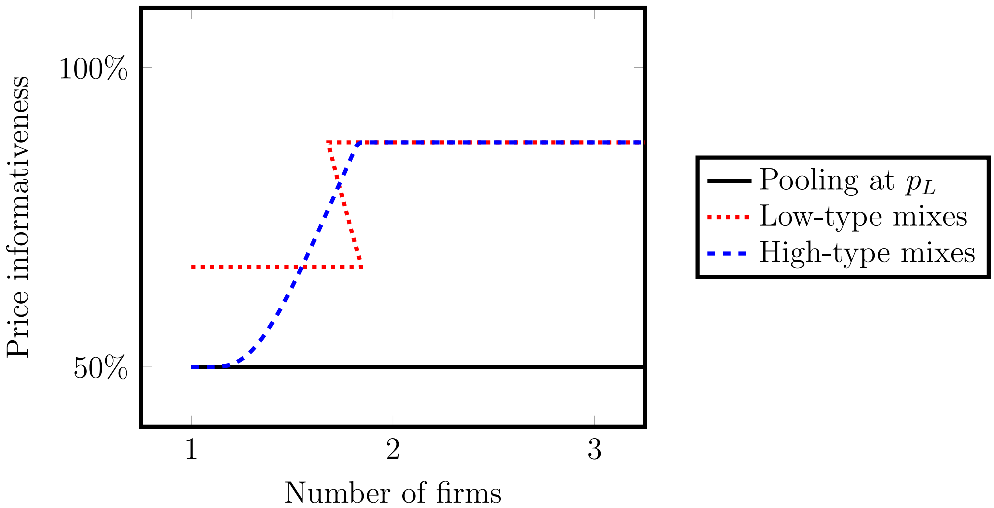

Figure 6: Price informativeness as competition increases

Notes: This figure is created with the parameters used in the experiment. There always exists an $N \in \mathbb{N}$ such that informativeness has fully converged for all $n \geq N$; this N is always greater than 1, but can be made arbitrarily large by the choice of parameters. For clarity, this figure graphs the equilibria as n varies continuously, even though, of course, only integer values of n make sense.

If buyers were to believe that, in a pooling scenario, high-quality firms would be much more likely to deviate to high prices than low-quality firms, then a pooling equilibrium could not be sustained. Instead, subjects would get to the informative equilibrium when $n = 1$ and prices would become more informative when n increases from 1 to 2 firms. On the other hand, as n increases, the set of others' potential strategies to which setting the low price is a firm's best response grows. So, in an environment with some strategic uncertainty or noise, we might expect the low price to be set more frequently by both types of firms as competition increases. The logit QRE, which generally selects from the set of sequential equilibria based on a broad notion of risk-dominance, correspondingly selects the informative equilibrium when $n = 1$, but then selects the pooling equilibrium when $n > 1$.

4 Experimental Design

Experiments were run in-person at the University of Virginia, with a sample of 156 undergraduate students over 14 sessions. In each session, subjects played the game 10 times per treatment. At the beginning of each treatment, subjects were chosen to be either buyers or sellers, and then each buyer was randomly matched with either one or two sellers and the game was played. In 7 sessions, there was only 1 seller per buyer, and in 7 sessions there were 2 sellers per buyer.[4]

In the baseline specification, the value of the high-quality product was $v_H = 200$ and the value of the low-quality product was $v_L = 100$. The firms were restricted to either a high price of $p_H = 160$ or a low price of $p_L = 80$. The per-unit costs of production were $c_H = 40$ for a high-quality product and $c_L = 0$ for a low-quality product. Subjects were paid for every decision, and real-money payoffs were scaled down to target \$30 per participant on average. The payment scale factors were fixed and told to participants in advance.

After the baseline specification, subjects played a specification where sellers were no longer restricted to two prices. In this specification, sellers were able to choose a price

[4]Initially, the design was within-subjects. Subjects played one treatment ($n = 1$ or $n = 2$) first, and then played the other afterward. But after finding significant order effects, I dropped all but the first treatment each session, to include only data where subjects do not have beliefs that are primed by earlier treatments. The Wilcoxon test for order effects rejects that treatment order is not a determinant of average prices with a p-value of 0.0052.

between 20 and 200 in increments of 20. This treatment examines the robustness of the results relative to a more general model where sellers can choose any price.

5 Results

Figure 7 plots the seller choice probabilities in each treatment. With a single seller, observations appear to be closer to the partially separating (informative) equilibrium, and further from the pooling (uninformative) equilibrum. In contrast, observations with two sellers have shifted closer to the uninformative pooling equilibrium. Buyer strategies and additional figures are given in the appendix.

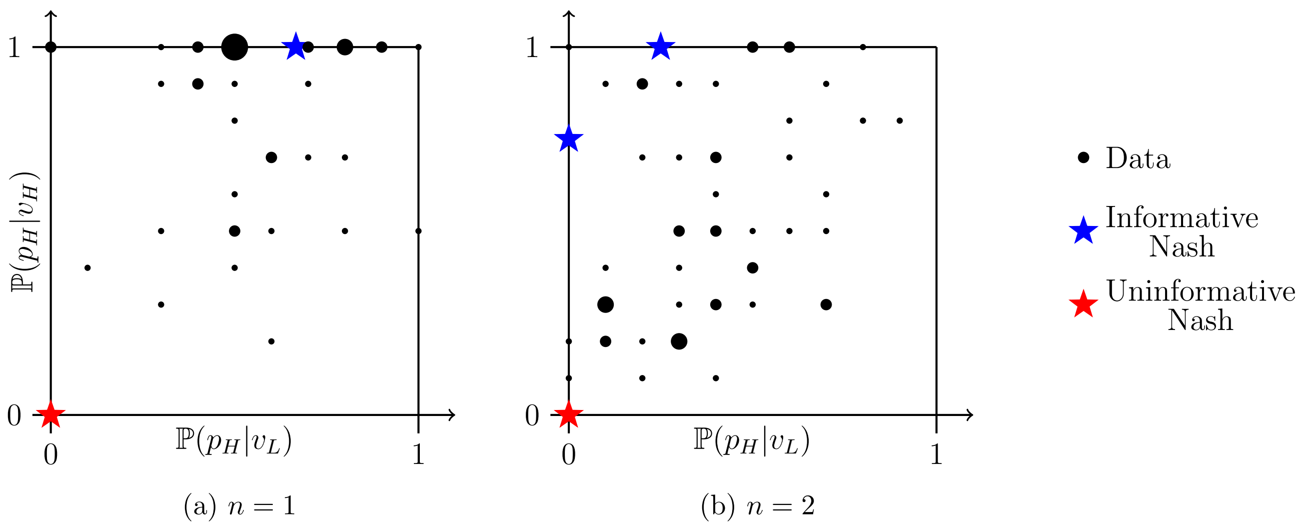

Figure 7: Empirical Seller Strategies

Notes: With a single seller, data is closer to the partially informative equilibrium. With two sellers, data moves towards the uninformative pooling equilibrium. Data is by-seller. The size of a datapoint represents the number of sellers who empirically played that strategy.

Table 1 reports the mean-squared error between the empirical choice probabilities and the choice probabilities in each Nash equilibrium. The partially separating equilibrium clearly minimizes the MSE when there is only one seller, but for two sellers, MSE is similar across all three Nash equilibria. Bootstrapping shows that, in every resampling of the $n = 1$ data, MSE selects the partially informative Nash, while resampling the $n = 2$ data leads to MSE selecting the pooling Nash about 47% of the time.

$n = 1$ firm			
	Low-types mix		**Pooling**
MSE	0.025 (0.006)		0.241 (0.030)
likelihood	1.000		0.000
$\mathbb{P}(p_H \vert v_H)$	1.000		0.000
$\mathbb{P}(p_H \vert v_L)$	0.667		0.000
informativeness	0.667		0.500
$n = 2$ firms			
	Low-types mix	**High-types mix**	**Pooling**
MSE	0.109 (0.015)	0.111 (0.014)	0.111 (0.024)
likelihood	0.484	0.050	0.466
$\mathbb{P}(p_H \vert v_H)$	1.000	0.750	0.000
$\mathbb{P}(p_H \vert v_L)$	0.250	0.000	0.000
informativeness	0.875	0.875	0.500

Table 1: Nash Equilibrium Selection using MSE

Notes: Standard errors are bootstrapped at the session level. The likelihood of an equilibrium is simply the proportion of times the resampled data led to a mean-squared error that selected that equilibrium.

Behavior in the experiment appears quite noisy: in most sessions, both high- and low-quality firms set both prices with significant probability. This contrasts with the Nash equilibria, in which at least one seller is always playing a pure pricing strategy. So, fitting the data to the Nash equilibria may not be realistic. Technically, no Nash is selected by the data since the likelihood of the data coming from any Nash is zero.

I fit a more realistic model by including noise in subject behavior. I've chosen to include noise using quantal response; an alternative parameterization using trembling hand (Selten 1975) is given in the appendix. While adding noise makes the model realistic, the intuition for equilibrium selection can change; adding noise fundamentally changes the equilibria and high levels of noise can lead to equilibria substantially different from any Nash.

5.1 Quantal Response

The quantal response model of agent behavior comes from McKelvey and Palfrey (1995, 1998). It is a generalization of the Nash equilibrium in which agents do not play their best responses with probability 1; instead, agents simply play "better responses" more frequently than "worse responses". This is achieved by assuming that agents experience some noise in the perceptions of their payoffs. Since agents have consistent beliefs, they understand that they and their fellow agents experience this noisy perception, and they respond to it.

The most common form of quantal response equilibrium is the logit quantal response equilibrium, in which the noise is assumed to come from a type-I extreme value distribution with precision λ.[5] When precision is zero, the noise overwhelms the true payoffs and agents choose uniformly randomly over their possible strategies. As precision tends to infinity, agents play their best response with probability approaching 1, and the quantal response equilibrium converges to a Nash equilibrium. I estimate the quantal response equilibria using the path-following procedure developed by Turocy (2005, 2010).

Holt, Goeree, and Palfrey (2016) note that the quantal response model is useful in two ways. The first is that adding noise can be necessary to create a non-degenerate likelihood which can then be used to estimate other parameters, and the second is that the noise itself may be an important feature of the data. I use QRE for both of these reasons. In the first case, I use noise to create a likelihood function with which I can estimate a finite mixture model to see how likely each equilibrium is to be selected. The noise allows me to decide which Nash is closer to the data even when the data may not perfectly align with either Nash.

But noise is also an important consideration in its own right. One big reason that prices may become less informative when competition increases is that competition increases the basin of attraction for setting the low price. That is, as competition increases, the space of others' strategies for which setting the low price is a best response gets larger. Of course, this is not Nash intuition. If agents are not noisy, the amount of others' strategies for which setting the low price is a best response is irrelevant; all that matters is whether setting the low price is a best response to others' particular equilibrium strategy. But, in a more

[5]Sometimes the quantal response is parametrized instead by the scale parameter of the logit errors, $\mu = 1/\lambda$. For the proofs, I follow Turocy (2005) in parametrizing the QRE by $\nu = \lambda/(1+\lambda)$.

complex world where agents are unsure of their opponents' actions because their opponents are noisy, the size of the basin of attraction is a determinant of the equilibrium selection.

This intuition is also associated with the "main branch" of the QRE correspondence. Out of potentially many quantal response equilibria, only one (almost always) constitutes a *continuous path* from infinite noise to zero noise. Since this is a homotopy from uniform randomization to a unique Nash equilibrium, it selects a Nash based on something like risk dominance–which Nash would be selected if agents started by thinking that all strategies would be played with equal probability and then updated smoothly from there.[6] In keeping with this intuition, Turocy (2005) has proved that the main branch of the QRE always selects the risk-dominant equilibrium in 2x2 games.[7] In this game, the main branch of the QRE correspondence selects the informative equilibrium when there is only one seller, but selects the pooling equilibrium where sellers always set the low price once there are two sellers.

Figure 8 shows the logit QRE correspondence for 1 and 2 sellers, as well as the equilibria at the likelihood-maximizing level of noise. Note that the logit QRE makes a strong prediction that the one-seller data *always* selects the branch that converges to the informative equilibrium. This is because, in the logit QRE dynamics, agents play better responses more often than worse responses. In a pooling equilibrium where sellers are almost always setting the low price, high-quality sellers still have a much greater benefit from deviating to the high price *relative to low-quality sellers*. Thus, in a logit QRE, high-quality sellers *do* deviate much more than low-quality sellers. So even if a deviation to the high price is quite unlikely, if it did occur, buyers would have to assume that it is most likely a high-quality product and would choose to buy. Thus, firms would benefit from deviating.

The best-fit QRE mirrors the data in that behavior selects the more informative equilibrium with one seller and the less informative equilibrium with two sellers. Also like the data, logit QRE involves so much noise that, while differences in average prices are large

[6]This is true if the main branch is monotonic in λ. If instead the main branch bends back on itself before ultimately converging to a Nash, then smoothly updating as λ increases is impossible, even though the main branch still selects a unique Nash.

[7]The notion of risk-dominance is itself not defined for more general games.

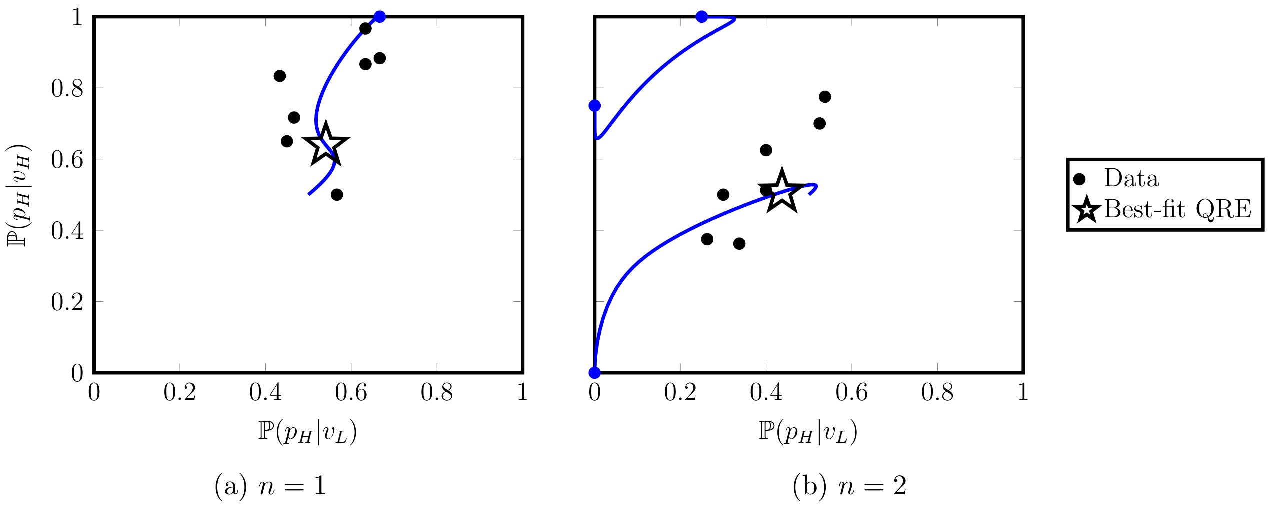

Figure 8: QRE Seller Strategies

Notes: With a single seller, there is a single QRE branch that converges to the informative equilibrium. With two sellers, there are multiple branches, but the data selects the main branch, which converges to the pooling equilibrium. The two-seller best-fit QRE is less informative than the one-seller best-fit QRE, although the difference is very slight since both equilibria involve significant noise. Data is by-session: session-level data must be used to fit the QRE since random matching means individual games in the same session are not independent.

between the one-seller and two-seller treatments, differences in informativeness are small. Prices are less informative with greater competition but the difference is very slight (about one percentage point in the best-fit QRE).

5.2 Non-parametric Tests

Parametric models are valuable because they can tell a story about deviations of behavior from theory, or pin down model parameters that underly the data. But the theory is also sensitive to how noisiness in behavior is implemented, and fitting a model like quantal response is a lot to ask of just a few observations. Thus, it is also valuable to see what can be proved from the data without any assumptions on the underlying behavioral model that subjects are following.

A non-parametric test cannot say anything about equilibrium selection (since it is agnostic about what an equilibrium is) but it can test hypotheses that depend purely on the data. Price informativeness itself is simply a feature of the data that can be calculated directly from the seller choice probabilities. Thus, the difference in informativeness between the two treatments can be tested using a non-parametric permutation test without any assumptions on the underlying data-generating process. Other differences, such as decreases in prices, can also be tested non-parametrically.

The null hypothesis of each permutation test is that the distribution of the data does not change with the number of sellers. Tests for price decreases are one-sided, since theory predicts that prices should weakly decrease, and the test for changes in informativeness is two-sided since theory is agnostic about the effect of treatment on informativeness. Each test permutes which data points are assigned to each treatment, and for each permutation, calculates the new difference in means across treatments. Under the null hypothesis, permuting which data is assigned to which treatment changes nothing, and thus the true difference in means from the actual data should not be too extreme relative to the differences in means created from the permutations of the data. Table 3 displays the results.

Effect of competition	Sample size per treatment	Number of permutations	More extreme permutations	p-value
High-quality price decreases	7	3432	43	0.0125
Low-quality price decreases	7	3432	31	0.0090
Average price decreases	7	3432	27	0.0079
Informativeness decreases	7	3432	352	0.1026

Figure 9: Permutation Tests

Notes: Each test calculates the average difference in the outcome variable between the two treatments (one vs. two sellers) for every possible way to permute which session belongs to which treatment. The p-value is simple the proportion of those permutations that yield differences more extreme than the actual data. Under the null hypothesis that the treatment does not affect the outcome, an actual test statistic that is quite extreme relative to the permutations is unlikely. The tests of price decreases are one-sided (since Nash theory predicts them to weakly decrease for my experimental parameters) and the test for informativeness is two-sided since theory is agnostic about the sign of the effect.

6 Conclusion

This paper examines the effect of competition in situations where consumers are sufficiently uninformed that they use prices as a guide to the quality of products; where prices convey some, but not all information about product quality. The model presented here is a simple and straightforward way that this situation can come about. The experiment presents suggestive evidence that the extent to which prices convey information about product quality decreases with competition. While competition clearly benefits consumers by lowering prices, this benefit may be tempered by a decrease in the *informational* value of prices.

Although the informational role of prices is difficult to assess in real-world markets, it could have serious welfare implications. Most obviously, more informative pricing may lead those consumers who value quality most highly to discern which are the best products, while less fastidious consumers settle for low-quality products at lower prices. Thus informative pricing can benefit welfare through increased allocative efficiency. If new firms entering a market cannot (or do not) credibly signal their quality, their entrance might decrease allocative efficiency, potentially leading high-quality firms to exit the market, even when the additional value to consumers of high-quality products exceeds their additional cost of production.

This model is intended as a simple example to show how competition might affect informativeness; there are many other ways that prices might convey information. Some channels (certifications, informed consumers, repeat customers, brand loyalty) quickly lead to perfect separation between firms, where consumers perfectly learn the quality prior to buying, and are never surprised. This may occur in many markets, but not the situations studied in this paper where ex-post regret occasionally occurs.

Nevertheless, there are many other models that do account for partial informativeness and occasional ex-post regret. One possibility is that consumers may be *initially* uninformed, but can become informed after incurring a cost of time or effort. Another possibility is that the space of product qualities is multi-dimensional, so that a (one-dimensional) price cannot possibly convey all the relevant quality information to consumers, who may differ in their relative value for different features of the product. These more involved models are left for future work.

Appendices

A Derivation of equilibria

The single-crossing result is shown in the paper, which implies there are only three potential types of equilibria: pooling at p_L, low-types mixing while high-types set p_H, and high-types mixing while low-types set p_L. All three turn out to be possible, although high-mixing equilibria only appear once n is sufficiently high.

First, I will calculate the probabilities of sale explicitly. Let R_H denote the probability with which a buyer will buy a high-priced product when only high-priced products are available, and let R denote the probability a buyer will buy a high-priced product when both high- and low-priced products are available. Then let x denote the probability that a seller sets the high price, unconditional on quality. Since each type of seller is equally likely, this means

$$x = \frac{1}{2}\left(\mathbb{P}(p_H|v_H) + \mathbb{P}(p_H|v_L)\right)$$

Then, the probability of an individual seller making a sale when setting the high price is

$$
\begin{aligned}
\mathbb{P}\left(\text{sale}|p_H\right) &= \mathbb{P}\left(\text{no others set } p_L\right)\frac{R_H}{n} + \sum_{i=1}^{n-1}\mathbb{P}\left(i \text{ others set } p_L\right)\frac{R}{n-i} \\
&= x^{n-1}\frac{R_H}{n} + \sum_{i=1}^{n-1}\binom{n-1}{i}x^{n-i-1}(1-x)^i\frac{R}{n-i} \\
&= x^{n-1}\frac{R_H}{n} + \left(\frac{R}{nx}\right)\sum_{i=1}^{n-1}\binom{n}{i}x^{n-i}(1-x)^i \\
&= x^{n-1}\frac{R_H}{n} + \left(\frac{R}{nx}\right)(1-x^n-(1-x)^n) \\
&= \frac{1}{n}\left[R_H x^{n-1} + R\left(\frac{1-x^n-(1-x)^n}{x}\right)\right]
\end{aligned}
$$

or, in the edge case where $x = 0$, $\mathbb{P}(\text{sale}|p_H) = R$. An analagous derivation shows that the probability of an individual seller making a sale when setting the low price is

$$
\begin{aligned}
\mathbb{P}\left(\text{sale}|p_L\right) &= \mathbb{P}\left(\text{no others set } p_H\right)\frac{1}{n} + \sum_{i=1}^{n-1}\mathbb{P}\left(i \text{ others set } p_H\right)\frac{1-R}{n-i} \\
&= \frac{1}{n}\left[(1-x)^{n-1} + (1-R)\left(\frac{1-x^n-(1-x)^n}{1-x}\right)\right]
\end{aligned}
$$

or, in the edge case where $x = 1$, $\mathbb{P}(\text{sale}|p_L) = 1 - R$.

Each equilibrium will involve seller optimality (which will be a constraint on seller profits, and thus on $\mathbb{P}(\text{sale}|p_H)$ and $\mathbb{P}(\text{sale}|p_L)$) as well as buyer optimality and consistent beliefs, which will require that R and R_H be the correct buyer choice probabilities given buyer beliefs, which in equilibrium relate to seller choice probabilities. Buyer optimality requires that

$$R_H = \begin{cases} 1 \;; \mathbb{P}(v_H|p_H) > \frac{p_H - v_L}{v_H - v_L} \equiv \underline{\mu} \\ \text{free} \in [0,1] \;;\; \mathbb{P}(v_H|p_H) = \underline{\mu} \\ 0 \;;\; \mathbb{P}(v_H|p_H) < \underline{\mu} \end{cases}$$

$$R = \begin{cases} 1 \;; \mathbb{P}(v_H|p_H) - \mathbb{P}(v_H|p_L) > \frac{p_H - p_L}{v_H - v_L} \equiv \bar{\mu} \\ \text{free} \in [0,1] \;;\; \mathbb{P}(v_H|p_H) - \mathbb{P}(v_H|p_L) = \bar{\mu} \\ 0 \;;\; \mathbb{P}(v_H|p_H) - \mathbb{P}(v_H|p_L) < \bar{\mu} \end{cases}$$

and belief consistency requires that

$$\mathbb{P}(v_H|p_H) = \frac{\mathbb{P}(p_H|v_H)}{\mathbb{P}(p_H|v_H) + \mathbb{P}(p_H|v_L)}$$
$$\mathbb{P}(v_H|p_L) = \frac{\mathbb{P}(p_L|v_H)}{\mathbb{P}(p_L|v_H) + \mathbb{P}(p_L|v_L)}$$

A.1 Pooling at p_L

A continuum of such equilibria always exist (for any n). Since p_H is never set, buyer beliefs about the expected quality of a high-priced product are free. All that is required is that sellers do not prefer to deviate to the high price. Since the high-quality sellers have the greatest incentive to set high prices, it suffices to show that high-quality sellers prefer setting low prices.

We need that the high-quality seller gains weakly higher profits when setting p_L than when setting p_H, as long as every other firm also sets p_L. If $n > 1$, we need

$$
\begin{aligned}
\mathbb{P}\left(\text{sale}|p_L\right)\left(p_L - c_H\right) &\geq \mathbb{P}\left(\text{sale}|p_H\right)\left(p_H - c_H\right) \\
\mathbb{P}\left(\text{sale}|p_H\right) &\leq \frac{p_L - c_H}{p_H - c_H}\mathbb{P}\left(\text{sale}|p_L\right) \equiv k_H \mathbb{P}\left(\text{sale}|p_L\right) \\
R &\leq \frac{k_H}{n}
\end{aligned}
$$

So this equilibrium can occur as long as R is sufficiently low, and R is always sufficiently low because

$$
\mathbb{P}(v_H|p_H) - \mathbb{P}(v_H|p_L) = \mathbb{P}(v_H|p_H) - \frac{1}{2} \leq \frac{1}{2} < \bar{\mu}
$$

If instead $n = 1$, the condition is that $R_H \leq k_H$, which can occur as long as $\mathbb{P}(v_H|p_H) < \underline{\mu}$, which is possible since p_H is off the equilibrium path.

A.2 Low-type mixes, High-type sets p_H

In this equilibrium we need the low-type seller to be indifferent between the two prices. Single-crossing will then guarantee that the high-type seller prefers to set the high price.

$$
\begin{aligned}
\mathbb{P}\left(\text{sale}|p_L\right)\left(p_L - c_L\right) &= \mathbb{P}\left(\text{sale}|p_H\right)\left(p_H - c_L\right) \\
\mathbb{P}\left(\text{sale}|p_H\right) &= \frac{p_L - c_L}{p_H - c_L}\mathbb{P}\left(\text{sale}|p_L\right) \equiv k_L \mathbb{P}\left(\text{sale}|p_L\right) \\
R_H x^{n-1} + R\left(\frac{1 - x^n - (1-x)^n}{x}\right) &= k_L\left[(1-x)^{n-1} + (1-R)\left(\frac{1 - x^n - (1-x)^n}{1-x}\right)\right]
\end{aligned}
$$

where $x = \frac{1}{2}(1 + \mathbb{P}(p_H|v_L))$.

Now consider the regions of the buyer belief space in figure 2. In region A, the buyer believes high-priced products are never worth buying. Since the buyer will not buy high-priced products, it is in sellers' best interests to set the low price, so there is no informative equilibrium here.

In region C, buyers always prefer high-priced products, and so all sellers want to set high prices. This cannot be an equilibrium because consumers do not have consistent beliefs. If all sellers are setting high prices, then high-priced products are just as likely to be low-quality as high-quality, so buyers should not want to buy them.

Consider the boundary between regions A and B. In this case, the buyer always prefers low-priced products ($R = 0$), but will occasionally buy a high-priced product (with probability R_H) when no low-priced products exist. To be at this boundary, we need $\mathbb{P}(v_H|p_H) = \underline{\mu}$, which implies that the low-type sets the high price with probability $\mathbb{P}(p_H|v_L) = 1/\underline{\mu} - 1$. Then the equation for low-type indifference pins down R_H:

$$R_H = \frac{k_L}{x^{n-1}}\left[(1-x)^{n-1} + \frac{1 - x^n - (1-x)^n}{1-x}\right] = \frac{k_L(1-x^n)}{x^{n-1}(1-x)}$$

For this equilibrium to exist, we need $R_H \in [0, 1]$, which implies a condition on the number of sellers. There is an informative equilibrium where the low-type mixes in the boundary between regions A and B as long as

$$n \leq 1 + \frac{ln\left(\frac{k_L}{1-(1-k_L)x}\right)}{ln(x)}$$

A similar derivation applies to the boundary between regions B and C. In this case, we have $R_H = 1$ and $x = 1/\bar{\mu} - 1$. Again, low-type seller indifference pins down the buyer choice probability:

$$R = \frac{k_L\left(\frac{1-x^n}{1-x}\right) - x^{n-1}}{\frac{1-(1-x)^{n-1}}{x} + k_L\frac{1-x^{n-1}}{1-x}}$$

and for $R \in [0,1]$, we need

$$n \geq 1 + \frac{ln\left(\frac{k_L}{1-(1-k_L)x}\right)}{ln(x)}$$

Lastly, there can be an equilibrium in the interior of region B. In this case, we have $R_H = 1$ and $R = 0$, so low-type seller indifference implies

$$x^{n-1}(1-x) = k_L(1-x^n)$$

and this pins down a unique $x \in [0,1)$ whenever such a solution exists, and a solution exists when n is intermediate between the cutoffs for the two boundaries of region B.

A.3 High-type mixes, Low-type sets p_L

In this type of equilibrium, $\mathbb{P}(v_H|p_H) = 1$ and

$$\mathbb{P}(v_H|p_L) = \frac{1-\mathbb{P}(p_H|v_H)}{2-\mathbb{P}(p_H|v_H)}$$

This type of equilibrium can exist in region B and in the boundary between regions B and C. In the boundary, we have $R_H = 1$ and

$$\mathbb{P}(v_H|p_H) - \mathbb{P}(v_H|p_L) = \bar{\mu} \implies \mathbb{P}(p_H|v_H) = \frac{2\bar{\mu}-1}{\bar{\mu}},\ x = \frac{2\bar{\mu}-1}{2\bar{\mu}}$$

and high-type indifference pins down the buyer mixing probability R. This equilibrium exists if

$$n \geq 1 + \frac{ln\left(\frac{k_H}{1-(1-k_H)x}\right)}{ln(x)}$$

In region B, $R_H = 1$ and $R = 0$, and then high-type indifference pins down a unique x:

$$x^{n-1}(1-x) = k_H(1-x^n)$$

and this equilibrium exists as long as

$$2 \leq n \leq 1 + \frac{ln\left(\frac{k_H}{1-(1-k_H)x}\right)}{ln(x)}$$

B Additional Figures

Figure 10 shows how the basin of attraction for setting the low price grows as competition increases. While the size of the basin doesn't matter to an agent who perfectly predicts others' actions in equilibrium, it may matter for actual players who may be noisy or face strategic uncertainty. The basins for the low-quality seller look very similar, except that the area in which p_L is a best response is larger.

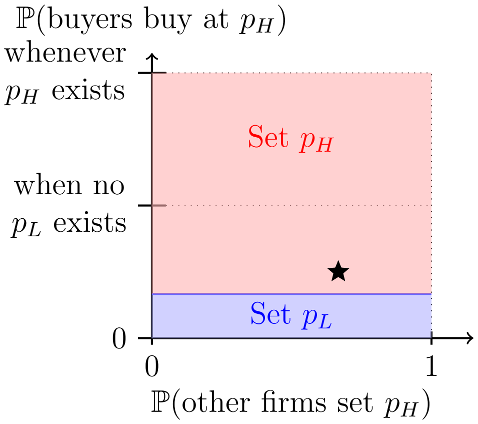

(a) $n = 1$

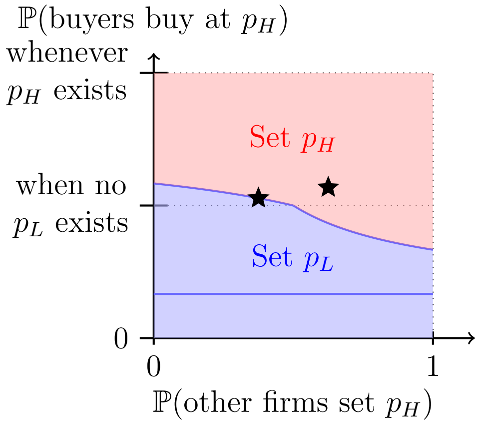

(b) $n = 2$

Figure 10: Basins of Attraction for High-quality Seller

Figure 11 gives the empirical buyer strategies. Movement from more separation when $n = 1$ towards more pooling when $n = 2$ is less clear here, partly because there are a continuum of pooling equilibria, which differ along one dimension of the buyer strategy.

Figure 12 shows the data aggregated by session. Aggregating by session is necessary for statistical tests since random matching within sessions means games played within a session are not statistically independent. Crosses show bootstrapped 90% confidence intervals.

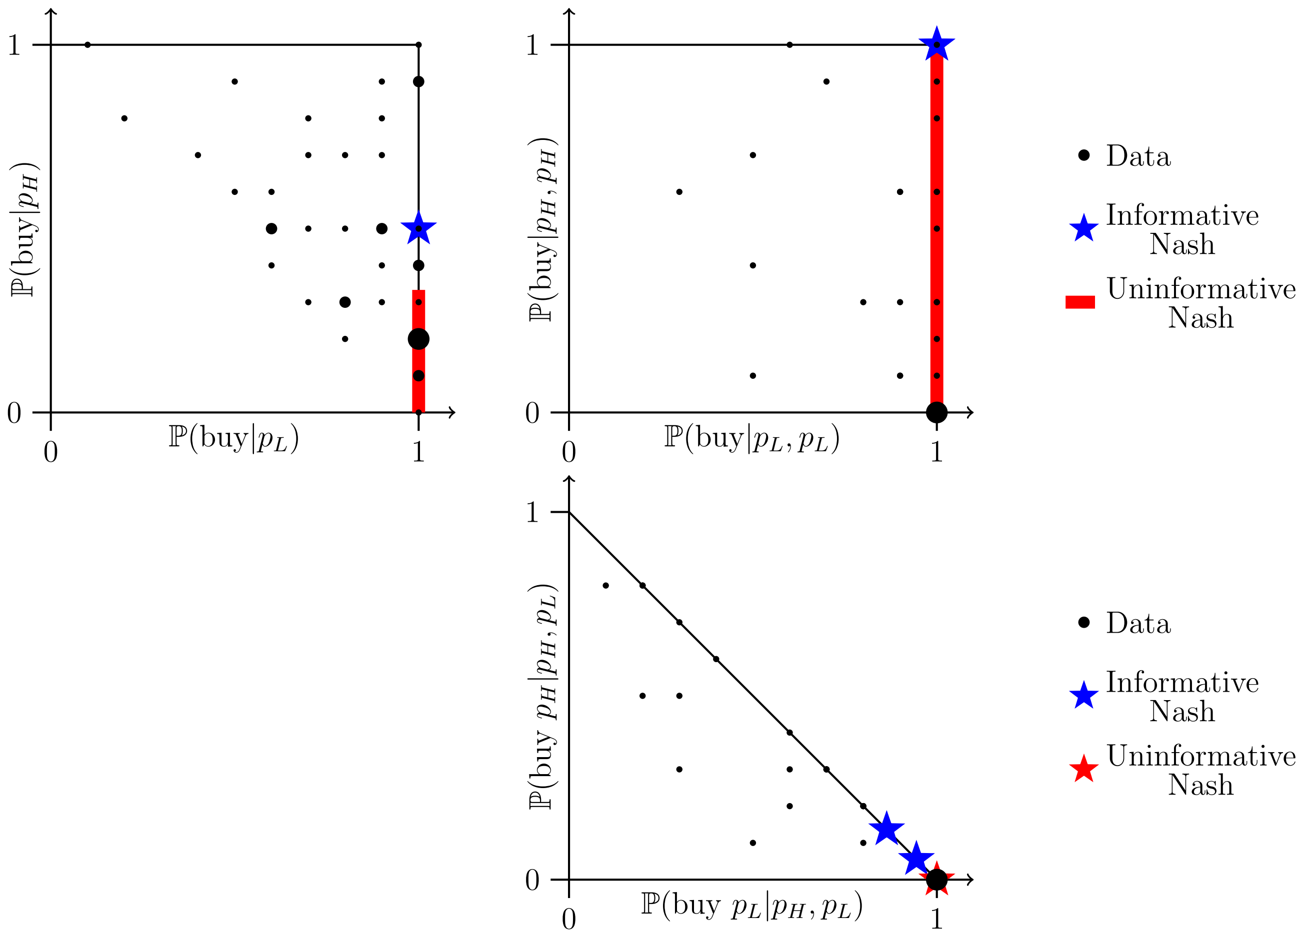

(a) $n = 1$

(b) $n = 2$

Figure 11: Empirical Buyer Strategies

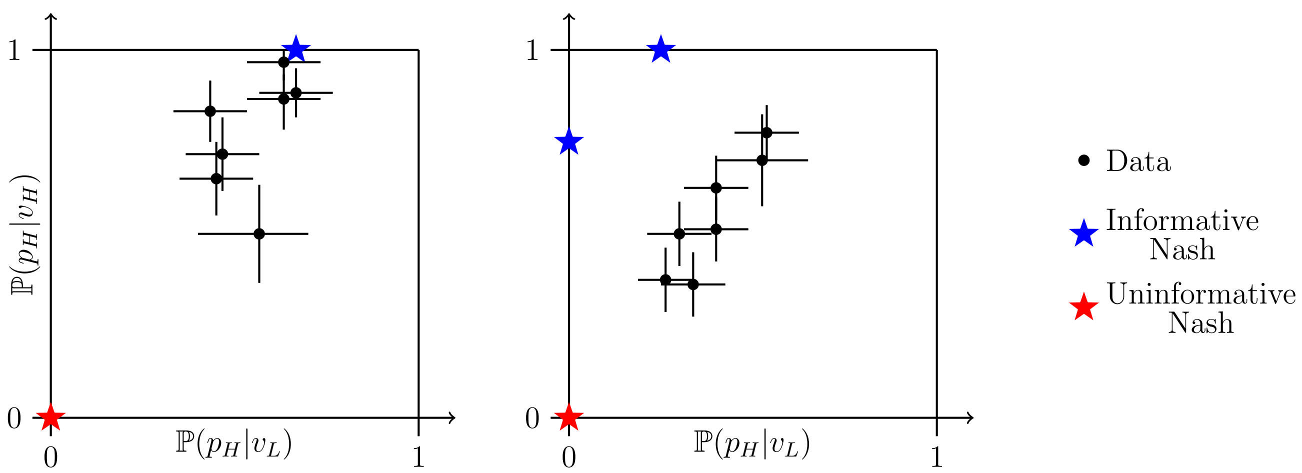

(a) $n = 1$

(b) $n = 2$

Figure 12: Empirical Seller Strategies by Session

C Instructions and zTree Code

All zTree code is available on github, at dvkwiat/compinf. Each treatment began with instructions which were shown on the screen and read aloud, followed by 10 rounds of decision screens. After each decision, subjects were given feedback on all information (qualities, prices, buying decisions, and payoffs). Then subjects were shown a history page to see the outcome of each round. Figures 13 and 14 show the seller and buyer decision screens in the 1-seller treatment.

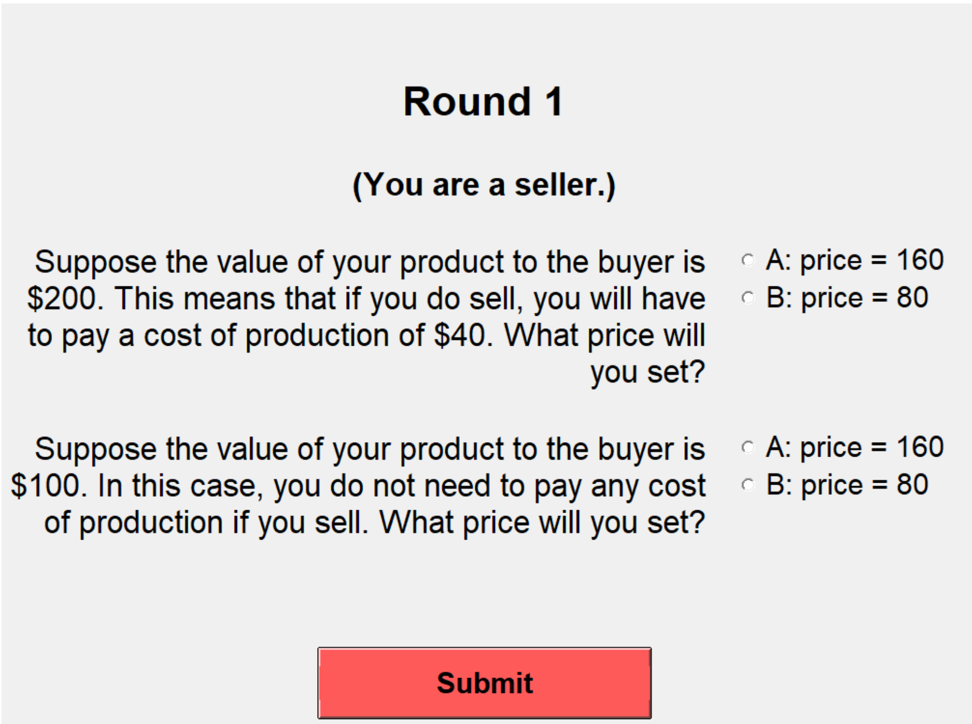

Figure 13: Seller Decision Screen

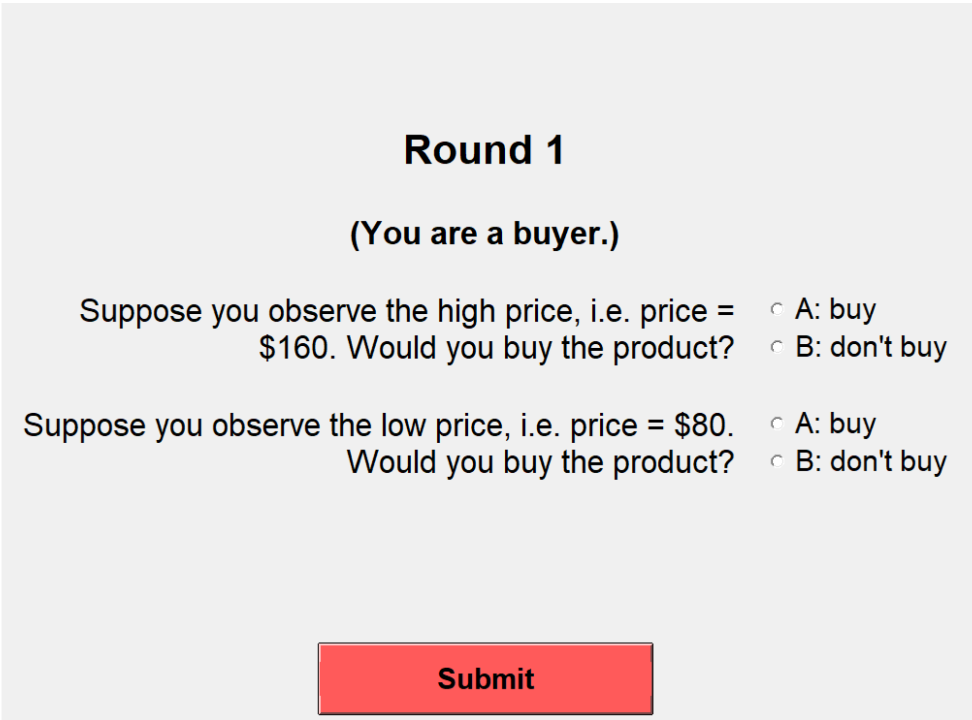

Figure 14: Buyer Decision Screen

D QRE Appendix

The quantal response equilibrium correspondence is found numerically using a path-following procedure outlined in Turocy (2005, 2010). When there is only one seller, there is a single QRE branch that converges to the unique informative equilibrium. Figures 15 and 16 show how QRE choice probabilities evolve as precision increases.

There are also a continuum of pooling Nash equilibria, but none are approached by a logit QRE. A proof of this follows: Consider a branch of the logit QRE correspondence parameterized by precision, λ. Let π_{bhh} and π_{bll} denote the buyer's probability of buying the product when it is priced high and buying the product when it is priced low, respectively. Let π_{shh} and π_{shl} denote the probability that a seller sets the high price when their product is high-quality and low-quality, respectively. At every point along the QRE branch, the following equations hold:

$$\begin{aligned}
\frac{\pi_{bhh}}{1-\pi_{bhh}} &= e^{\lambda(v_L - p_H + \mu_H(v_H - v_L))} \\
\frac{\pi_{bll}}{1-\pi_{bll}} &= e^{\lambda(v_L - p_L + \mu_L(v_H - v_L))} \\
\frac{\pi_{shh}}{1-\pi_{shh}} &= e^{\lambda((p_H - c_H)\pi_{bhh} - (p_L - c_H)\pi_{bll})} \\
\frac{\pi_{shl}}{1-\pi_{shl}} &= e^{\lambda((p_H - c_L)\pi_{bhh} - (p_L - c_L)\pi_{bll})}
\end{aligned}$$

where

$$\mu_H = \frac{\pi_{shh}}{\pi_{shh} + \pi_{shl}}, \ \mu_L = \frac{1-\pi_{shh}}{2 - \pi_{shh} - \pi_{shl}}$$

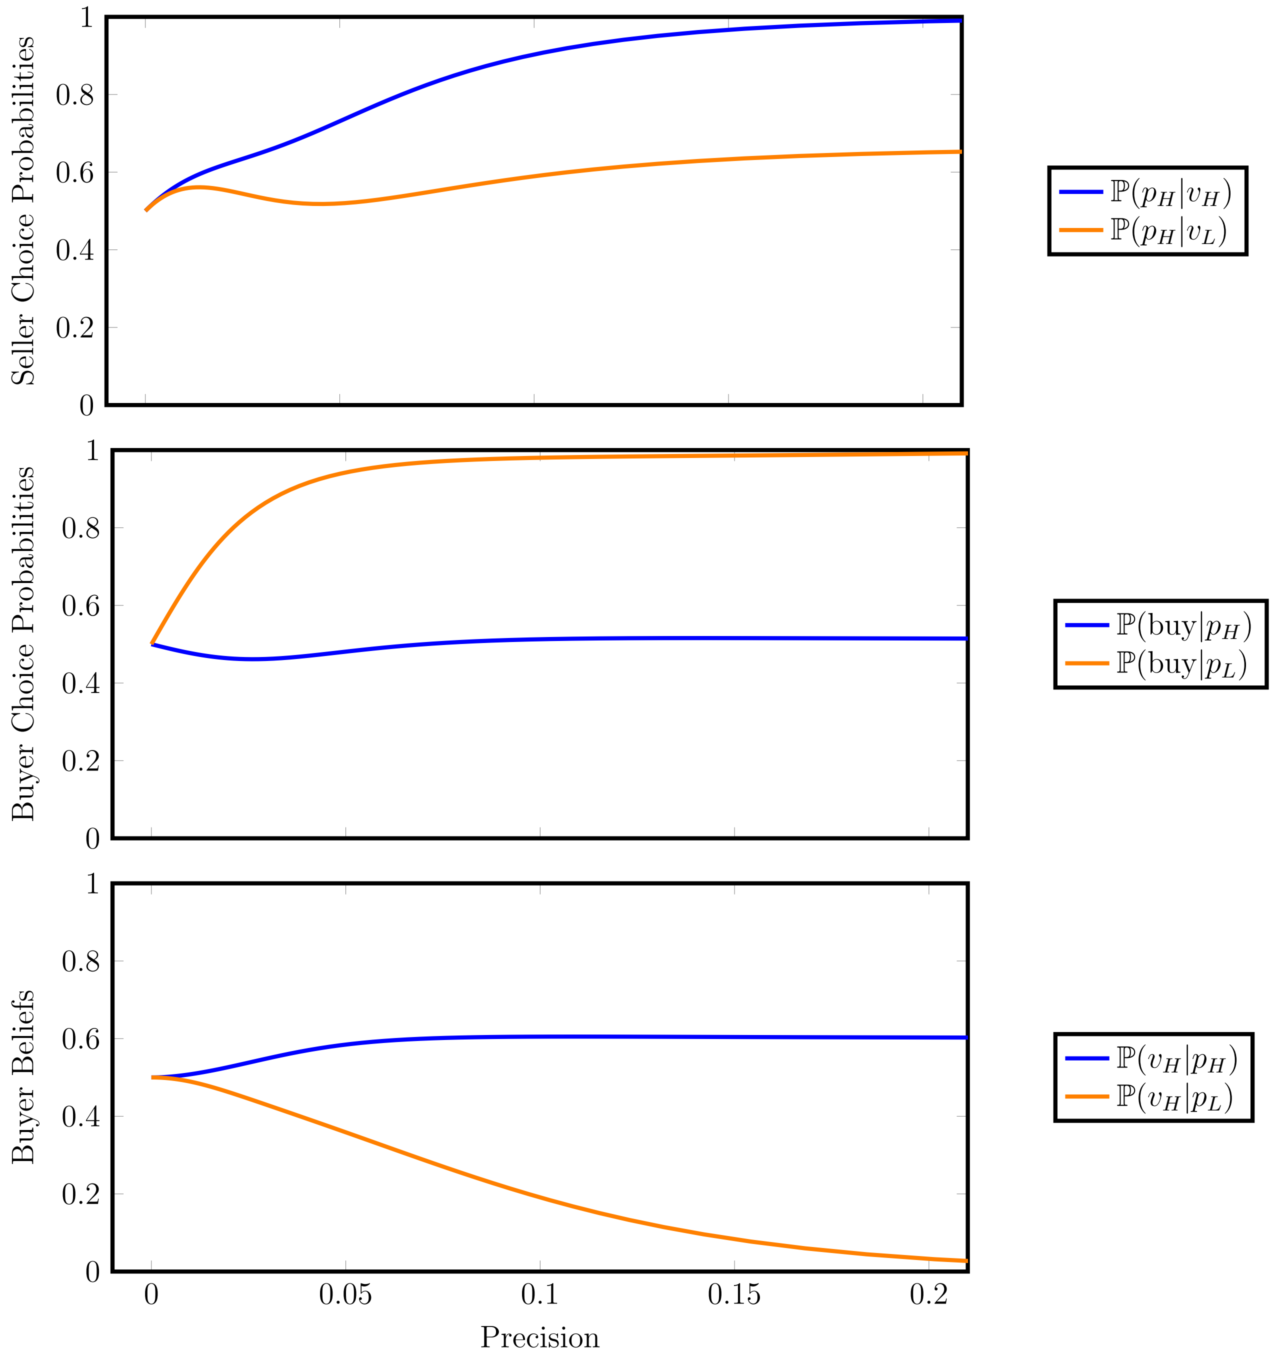

Figure 15: QRE correspondence for $n = 1$

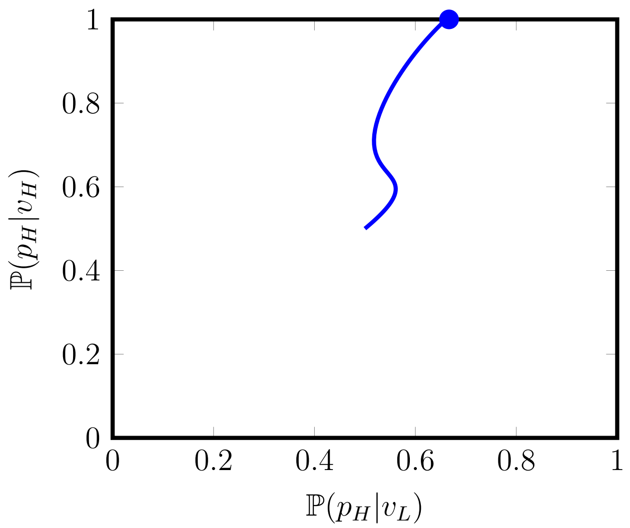

Figure 16: QRE Seller Strategies for $n = 1$

Taking logs and substituting $\nu = \lambda/(1+\lambda)$ yields

$$(1-\nu)\left(ln\left(\pi_{bhh}\right) - ln\left(1-\pi_{bhh}\right)\right) = \nu\left(v_L - p_H + \mu_H(v_H - v_L)\right) \tag{1}$$

$$(1-\nu)\left(ln\left(\pi_{bll}\right) - ln\left(1-\pi_{bll}\right)\right) = \nu\left(v_L - p_L + \mu_L(v_H - v_L)\right) \tag{2}$$

$$(1-\nu)\left(ln\left(\pi_{shh}\right) - ln\left(1-\pi_{shh}\right)\right) = \nu\left((p_H - c_H)\pi_{bhh} - (p_L - c_H)\pi_{bll}\right) \tag{3}$$

$$(1-\nu)\left(ln\left(\pi_{shl}\right) - ln\left(1-\pi_{shl}\right)\right) = \nu\left((p_H - c_L)\pi_{bhh} - (p_L - c_L)\pi_{bll}\right) \tag{4}$$

Suppose, for a contradiction, that the QRE branch converges to a pooling Nash equilibrium as $\lambda \to \infty$ (and thus $\nu \to 1$). This means that $\pi_{shh} \to 0$ and $\pi_{shl} \to 0$. First, notice that $\mu_L \to \frac{1}{2}$. Then from 2, we have that

$$(1-\nu)\left(ln\left(\pi_{bll}\right) - ln\left(1-\pi_{bll}\right)\right) \to \frac{v_L + v_H}{2} - p_L > 0$$

And since $\nu \to 1$, this implies that

$$ln\left(\pi_{bll}\right) - ln\left(1-\pi_{bll}\right) \to \infty \implies \pi_{bll} \to 1$$

By assumption, choice probabilities converge as precision approaches infinity, so let $\bar{\pi}_{bhh}$ be the limit of π_{bhh}. From 3, we have that

$$(1-\nu)\left(ln\left(\pi_{shh}\right)-ln\left(1-\pi_{shh}\right)\right)\rightarrow(p_H-c_H)\bar{\pi}_{bhh}-(p_L-c_H)$$

But note that the left-hand-side of this expression is eventually always less than or equal to zero, since by assumption, $\pi_{shh}\rightarrow 0$. So it must be that

$$(p_H-c_H)\bar{\pi}_{bhh}-(p_L-c_H)\leq 0 \implies \bar{\pi}_{bhh}\leq\frac{p_L-c_H}{p_H-c_H} \tag{5}$$

Now, substituting 3 and 4 into μ_H, we have that

$$\frac{1}{\mu_H}-1=\frac{1+e_H}{1+e_L}=\frac{1}{1+e_L}+\frac{1}{\frac{1}{e_H}+\frac{e_L}{e_H}}$$

where

$$e_H\equiv e^{\frac{\nu}{1-\nu}((p_L-c_H)\pi_{bll}-(p_H-c_H)\pi_{bhh})}$$
$$e_L\equiv e^{\frac{\nu}{1-\nu}((p_L-c_L)\pi_{bll}-(p_H-c_L)\pi_{bhh})}$$

Since $\bar{\pi}_{bhh}\leq(p_L-c_H)/(p_H-c_H)$ and $\pi_{bll}\rightarrow 1$, we know that $e_L\rightarrow\infty$, $e_H\not\rightarrow 0$, and

$$\frac{e_L}{e_H}=e^{\left(\frac{\nu}{1-\nu}\right)(c_H-c_L)(\pi_{bll}-\pi_{bhh})}\rightarrow\infty$$

Together, these imply that $\mu_H\rightarrow 1$. Then, from 1, we have that

$$(1-\nu)\left(ln\left(\pi_{bhh}\right)-ln\left(1-\pi_{bhh}\right)\right)\rightarrow v_H-p_H>0$$

which implies that

$$ln\,(\pi_{bhh}) - ln\,(1 - \pi_{bhh}) \rightarrow \infty \implies \pi_{bhh} \rightarrow 1$$

and this is a contradiction with 5 above. Therefore, there are no QRE branches that converge to pooling Nash equilibria.

Intuitively, the QRE is sidestepping the issue of off-equilibrium path beliefs. For every finite precision, all players play all their possible strategies with some positive probability, and thus beliefs are uniquely pinned down by rational expectations. As agents become very precise, high-type sellers have a greater incentive than low-type sellers to set the high price. For both sellers, that incentive must be vanishing, to sustain sellers never setting the high price in the limit. But for any finite precision, high-type sellers will set the high price relatively more than low-type sellers. As precision tends to infinity, both sellers set the high price less and less, but the comparative difference increases; the high-type seller is more and more likely to set the high price *relative to the low-type seller.* This means that consumers know, with probability approaching 1, that a high price must have come from a high-type seller. If consumers know this, they should converge to always buying when presented with a high price product. But this makes sellers regret their decision to set the high price with vanishing probability.

When there are two sellers, there are (for some levels of precision) three QREa. Two converge to the two informative Nash equilibria, and one (the main branch) converges to a pooling Nash equilibrium. Figures 17 and 18 show the choice probabilities in the QRE correspondence as precision increases.

Again, there are a continuum of pooling Nash equilibria, and only one is approached by a QRE. A proof of this follows. Consider a branch of the QRE correspondence that converges to a pooling Nash equilibrium, and suppose there are two or more sellers. Let π_{bhh} be the probability with which the buyer buys a product when only high-priced products are

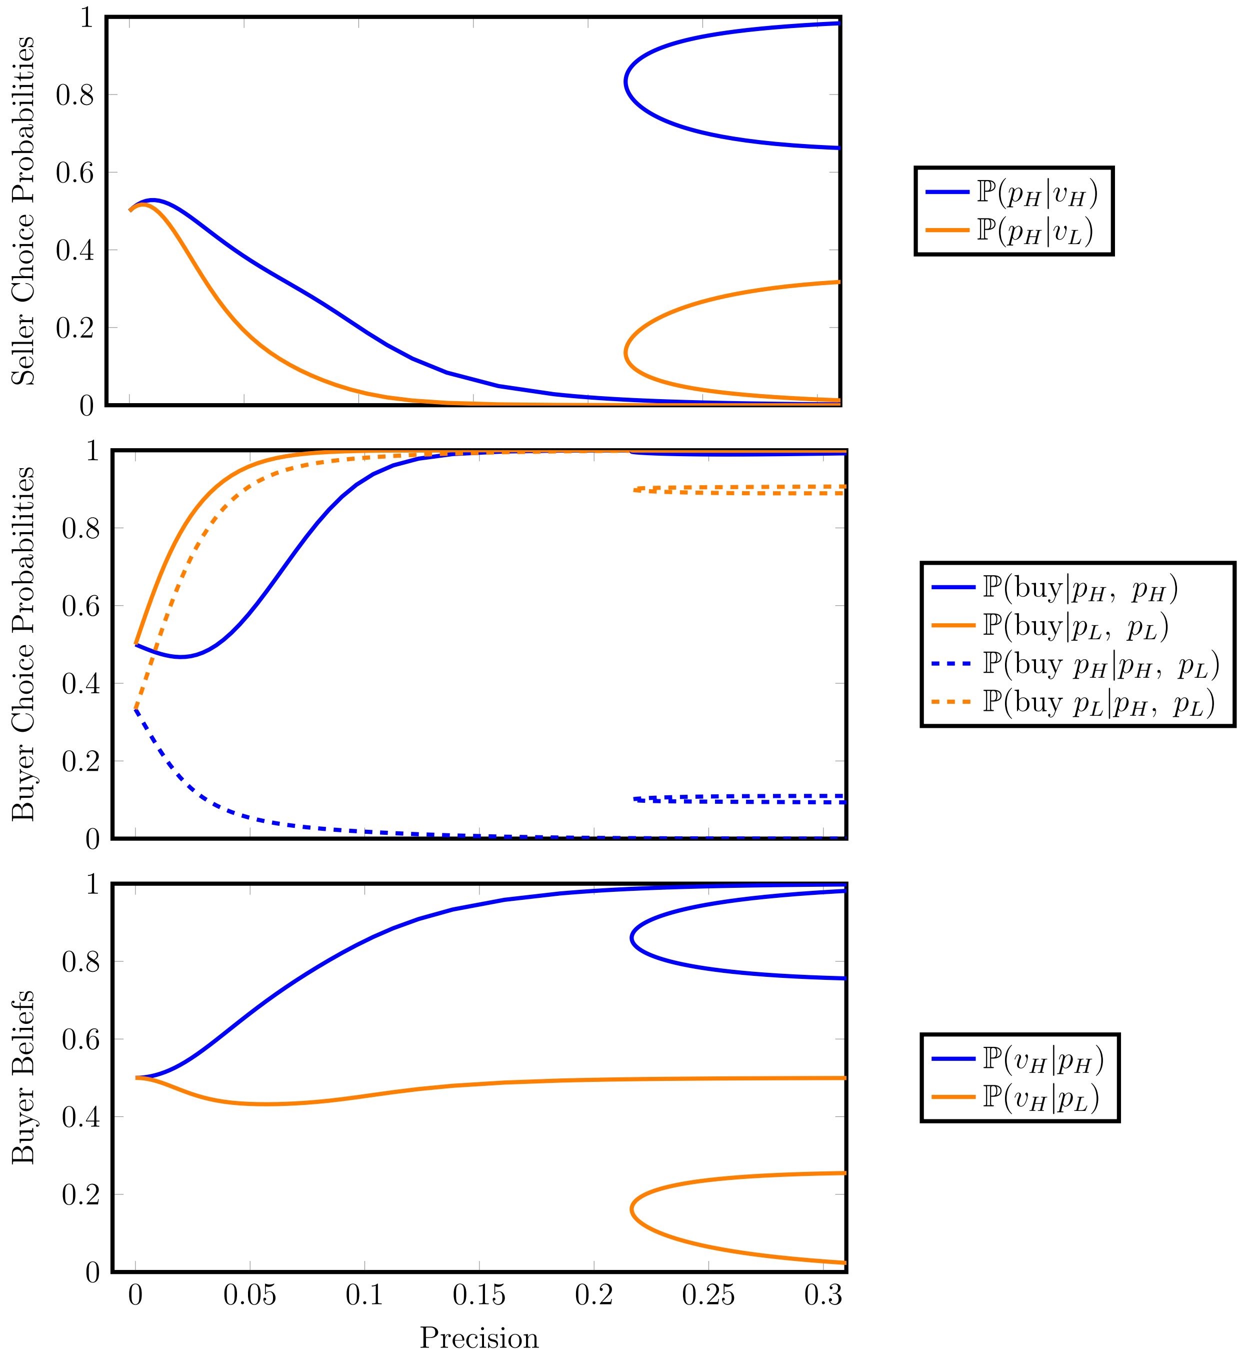

Figure 17: QRE correspondence for $n = 2$

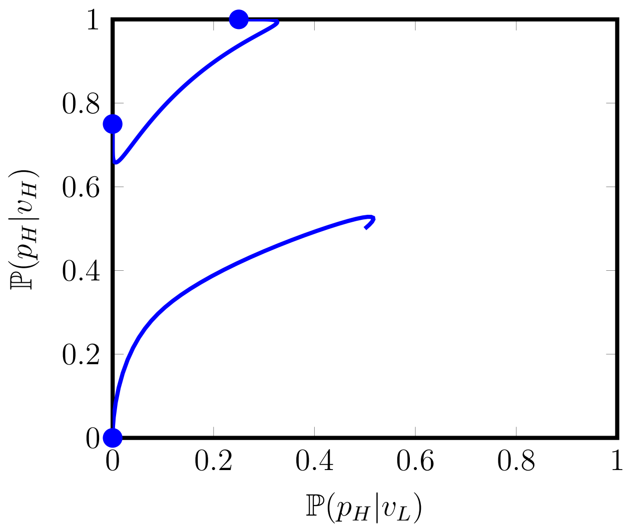

Figure 18: QRE Seller Strategies for $n = 2$

available, and let π_{bll} be the probability with which the buyer buys a product if all products are low-priced. Let π_{bhb} and π_{blb} denote the probability that a buyer buys a high-priced product or a low-priced product, respectively, when both are available. Define π_{shh} and π_{shl} as above, and again let $\nu \equiv \lambda/(1+\lambda)$. Every point on the QRE branch satisfies the following equations.

$$(1-\nu)\left(ln\left(\pi_{bhh}\right) - ln\left(1-\pi_{bhh}\right)\right) = \nu\left(v_L - p_H + \mu_H(v_H - v_L)\right) \tag{6}$$

$$(1-\nu)\left(ln\left(\pi_{bll}\right) - ln\left(1-\pi_{bll}\right)\right) = \nu\left(v_L - p_L + \mu_L(v_H - v_L)\right) \tag{7}$$

$$(1-\nu)\left(ln\left(\pi_{bhb}\right) - ln\left(1-\pi_{bhb}-\pi_{blb}\right)\right) = \nu\left(v_L - p_H + \mu_H(v_H - v_L)\right) \tag{8}$$

$$(1-\nu)\left(ln\left(\pi_{blb}\right) - ln\left(1-\pi_{bhb}-\pi_{blb}\right)\right) = \nu\left(v_L - p_L + \mu_L(v_H - v_L)\right) \tag{9}$$

$$(1-\nu)\left(ln\left(\pi_{shh}\right) - ln\left(1-\pi_{shh}\right)\right) = \nu\left((p_H - c_H)\mathbb{P}_H - (p_L - c_H)\mathbb{P}_L\right) \tag{10}$$

$$(1-\nu)\left(ln\left(\pi_{shl}\right) - ln\left(1-\pi_{shl}\right)\right) = \nu\left((p_H - c_L)\mathbb{P}_H - (p_L - c_L)\mathbb{P}_L\right) \tag{11}$$

where

$$\mathbb{P}_H = \frac{1}{n}\left[(\pi_{bhh} - \pi_{bhb})x^{n-1} + \pi_{bhb}\frac{1-(1-x)^n}{x}\right]$$
$$\mathbb{P}_L = \frac{1}{n}\left[(\pi_{bll} - \pi_{blb})(1-x)^{n-1} + \pi_{blb}\frac{1-x^n}{1-x}\right]$$
$$\mu_H = \frac{\pi_{shh}}{\pi_{shh}+\pi_{shl}},\ \mu_L = \frac{1-\pi_{shh}}{2-\pi_{shh}-\pi_{shl}},\ x = \frac{\pi_{shh}+\pi_{shl}}{2}$$

Because, by assumption, this branch converges to a pooling Nash equilibrium as $\nu \to 1$, we have that $\pi_{shh} \to 0$ and $\pi_{shl} \to 0$, as before. Thus, $x \to 0$ and $\mu_L \to 1/2$. Because $\mu_L \to 1/2$, we have (from 7 and 9) that

$$(1-\nu)\left(ln\left(\pi_{bll}\right) - ln\left(1-\pi_{bll}\right)\right) \to \frac{v_L+v_H}{2} - p_L \implies \pi_{bll} \to 1$$
$$(1-\nu)\left(ln\left(\pi_{bhb}\right) - ln\left(1-\pi_{bhb}-\pi_{blb}\right)\right) \to \frac{v_L+v_H}{2} - p_L \implies \pi_{blb}+\pi_{bhb} \to 1$$

Combining 8 and 9 yields

$$\frac{\pi_{bhb}}{\pi_{blb}} = e^{\left(\frac{\nu}{1-\nu}\right)(v_H-v_L)\left(\mu_H-\mu_L-\frac{p_H-p_L}{v_H-v_L}\right)}$$

Since $(p_H-p_L)/(v_H-v_L) > 1/2$ and $\mu_L \to 1/2$, the quantity

$$\mu_H - \mu_L - \frac{p_H-p_L}{v_H-v_L}$$

is eventually negative. Thus, as $\nu \to 1$, $\pi_{bhb} \to 0$ and thus $\pi_{blb} \to 1$. As in the proof for one seller above, we can use 10 and 11 to write

$$\frac{1}{\mu_H} - 1 = \frac{1+e_H}{1+e_L} = \frac{1}{1+e_L} + \frac{1}{\frac{1}{e_H}+\frac{e_L}{e_H}}$$

where

$$e_H = e^{\left(\frac{\nu}{1-\nu}\right)[(p_L - c_H)\mathbb{P}_L - (p_H - c_H)\mathbb{P}_H]}$$
$$e_L = e^{\left(\frac{\nu}{1-\nu}\right)[(p_L - c_L)\mathbb{P}_L - (p_H - c_L)\mathbb{P}_H]}$$

Since $x \to 0$, $\pi_{bhb} \to 0$, and $\pi_{blb} \to 1$, we know that $\mathbb{P}_L \to 1/n$ and $\mathbb{P}_H \to 0$, so

$$e_H \to \infty,\ e_L \to \infty,\ \frac{e_L}{e_H} = e^{\left(\frac{\nu}{1-\nu}\right)(c_H - c_L)(\mathbb{P}_L - \mathbb{P}_H)} \to \infty$$

So

$$\frac{1}{\mu_H} - 1 \to 0 \implies \mu_H \to 1$$

Lastly, if $\mu_H \to 1$, then by 6,

$$(1-\nu)\left(ln\left(\pi_{bhh}\right) - ln\left(1 - \pi_{bhh}\right)\right) \to v_H - p_H > 0 \implies \pi_{bhh} \to 1$$

Thus, only one pooling Nash equilibrium is approached by a QRE, specifically the Nash equilibrium where

$$\pi_{bhh} = 1,\ \pi_{bll} = 1,\ \pi_{bhb} = 0,\ \pi_{blb} = 1$$
$$\pi_{shh} = 0,\ \pi_{shl} = 0,\ \mu_H = 1,\ \mu_L = \frac{1}{2}$$

E Trembling-Hand Appendix

Here I have added noise using the trembling-hand model (Selten 1975). This model is paramaterized with a level of noise, ϵ. Agents play their best response with probability $1-\epsilon$, and with probability ϵ, they uniformly randomize over their available actions. As $\epsilon \to 1$, agent behavior is entirely random noise, and as $\epsilon \to 0$, the set of trembling-hand equilibria converge to a subset of the sequential Nash equilibria without noise.

When noise is implemented using trembling-hand, it often turns out to not effect the Nash equilibria in the limit as noise approaches zero. This holds true here as well. The logit quantal response model ruled out one of the sequential equilibria since the off-equilibrium-path beliefs were inconsistent with the logit dynamics, even before the QRE was fit to the data. With trembling-hand, all types of sequential equilibria are represented as limits of the trembling-hand equilibrium correspondence as noise approaches zero.

The trembling-hand choice probabilities are graphed in figures 19 and 20. The horizontal axis is $1 - \epsilon$ to be consistent with the QRE graphs above, where uniform randomization is on the left side of the graphs and convergence to Nash occurs on the right side.

As with quantal response, I use maximum likelihood to estimate the noise parameter of the trembling-hand model and the probabilities of each equilibrium being selected. The trembling-hand correspondence, along with the data and the best-fit equilibrium is given in figure 21.

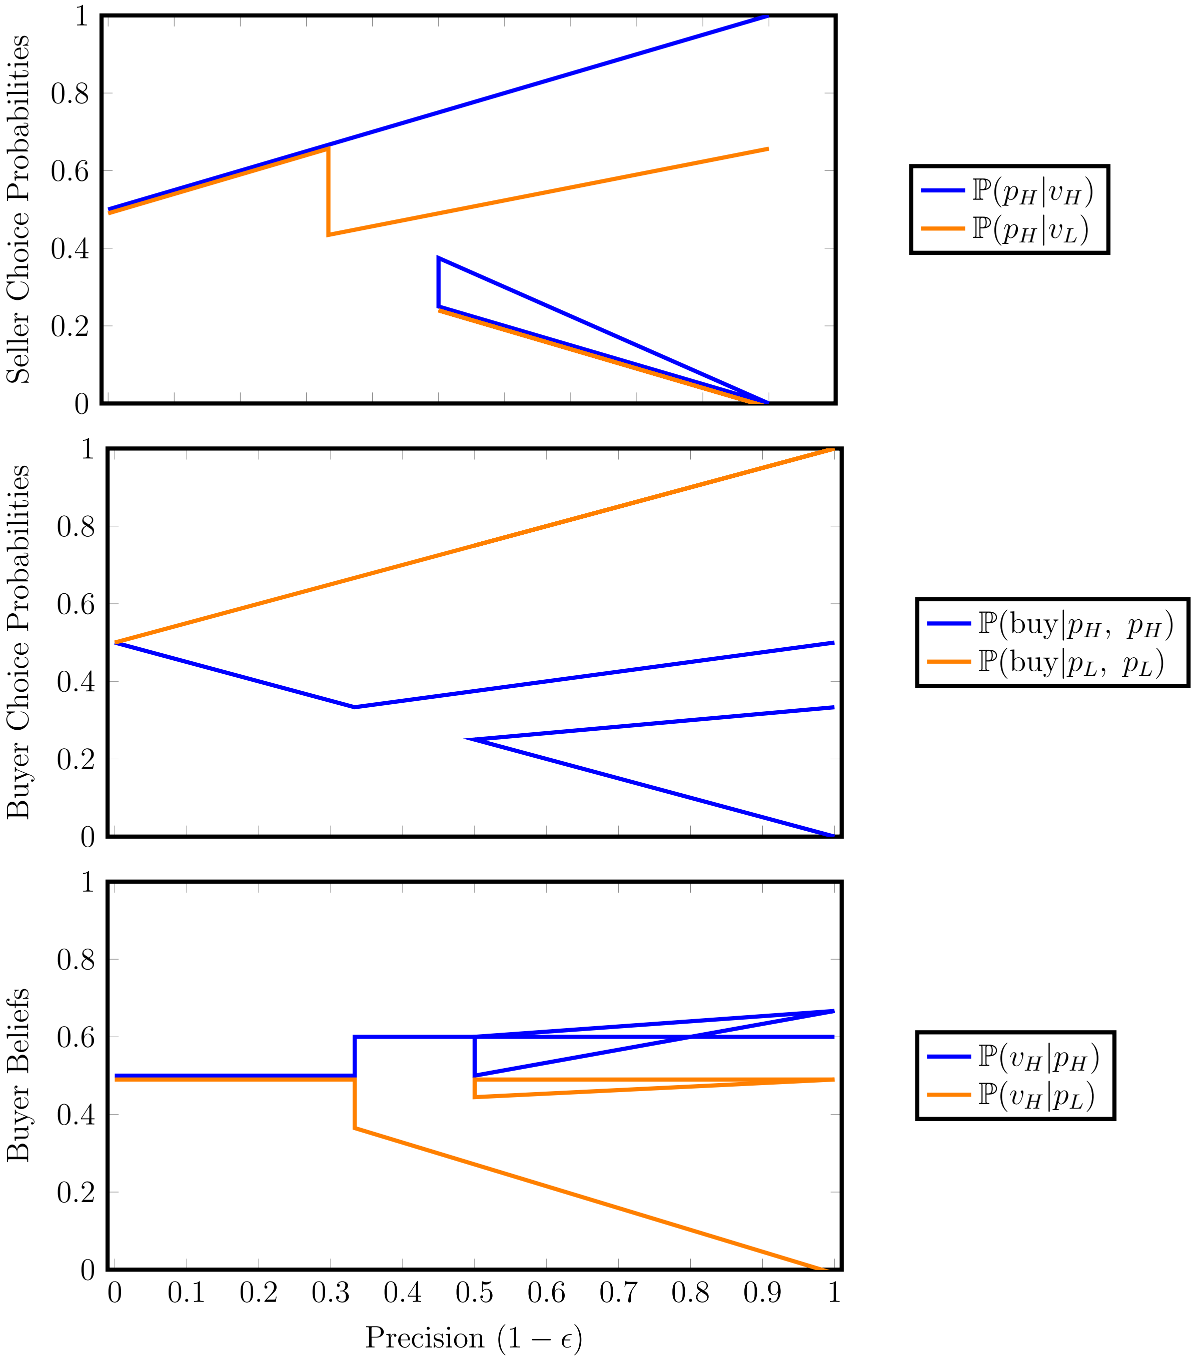

Figure 19: Trembling-hand correspondence for $n = 1$

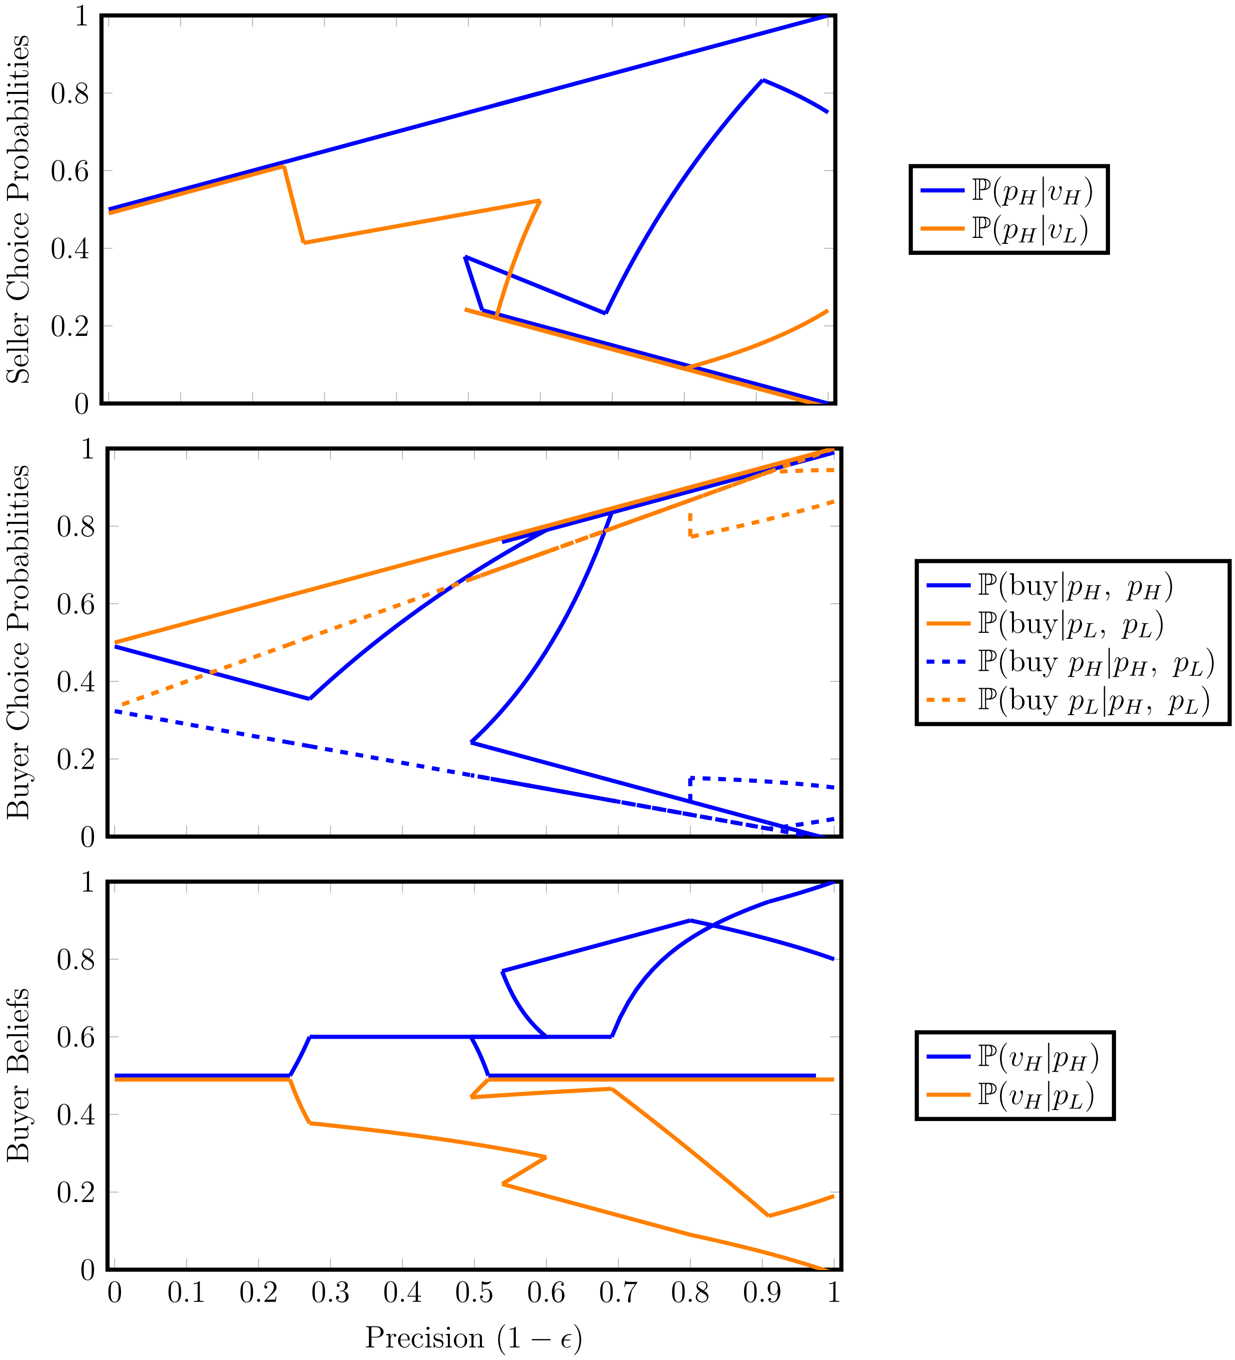

Figure 20: Trembling-hand correspondence for $n = 2$

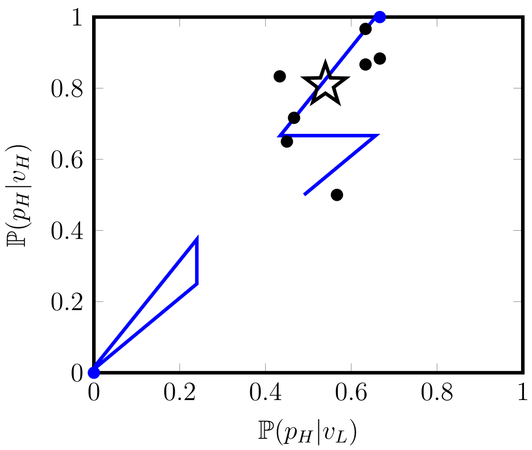

(a) $n = 1$

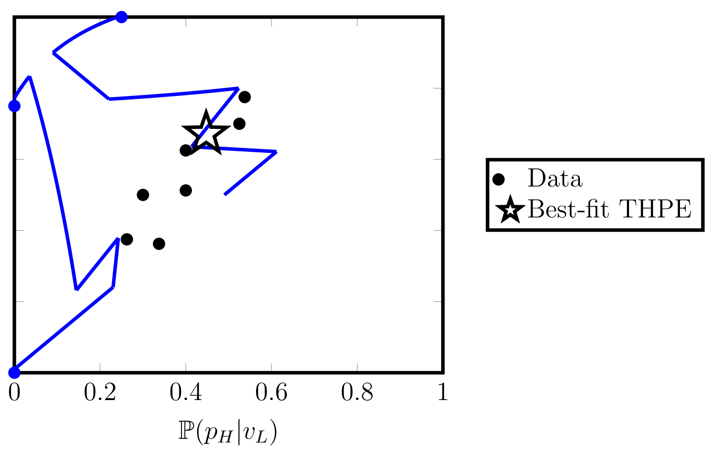

(b) $n = 2$

Figure 21: Trembling-Hand Seller Strategies

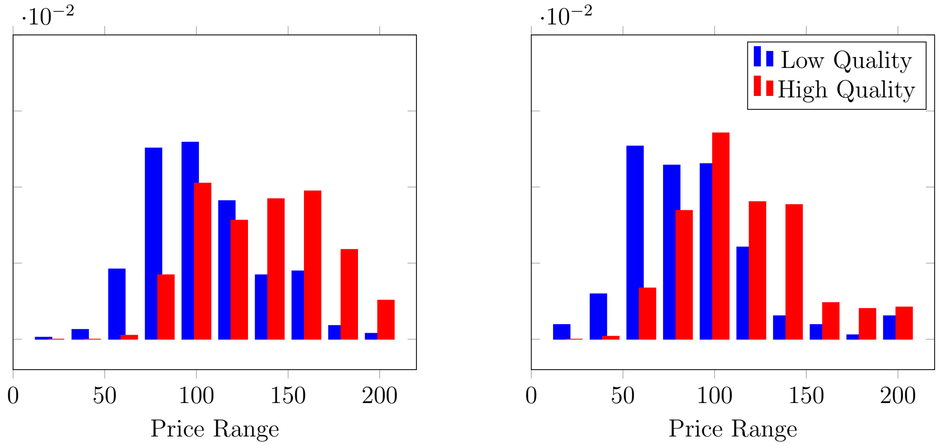

(a) $n = 1$ (b) $n = 2$

Figure 22: Histogram of Low and High Quality Prices

F Price Grid

The robustness treatment with the price grid had similar results. More competition caused lower prices and a slight decrease in informativeness, from 67.4% with a single seller to 66.1% with two sellers. The histograms of prices under the two treatments are given in figure 22.

G Results Including Treatments Run Second

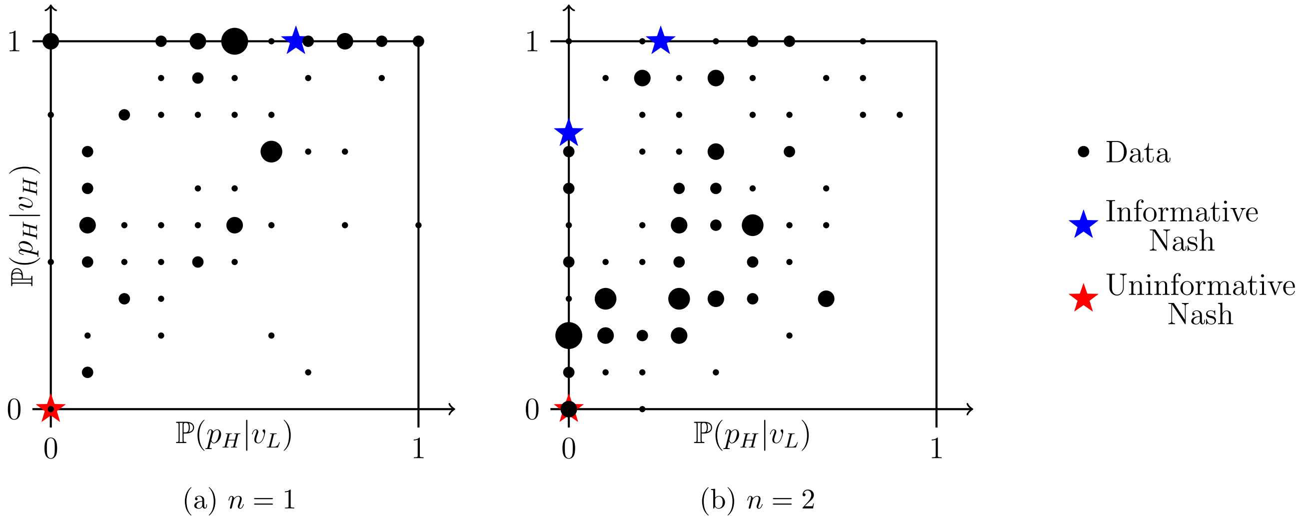

Figure 23: Empirical Seller Strategies

Figure 23 shows the empirical seller strategies for the one-seller and two-seller treatments. This graph includes sessions where the one seller treatment was following by the two-seller treatment (within the same subjects), and sessions where the two-seller treatment was followed by the one-seller treatment. As expected, beliefs are primed by whichever treatment occurred first, so there is persistence in the type of equilibria selected throughout treatments in a given session. This makes behavior appear more noisy, and can generally erode significance. Despite strong order effects, qualitative differences between the two treatments remain visible. Figure 24 shows the empirical buyer strategies including both treatment orders. In figure 25, data is aggregated by session.

The p-value for a decrease in informativeness is smaller when treatments run second are included (about 5.3 %); the beneficial effect of additional data outweighs the noisiness of that data.

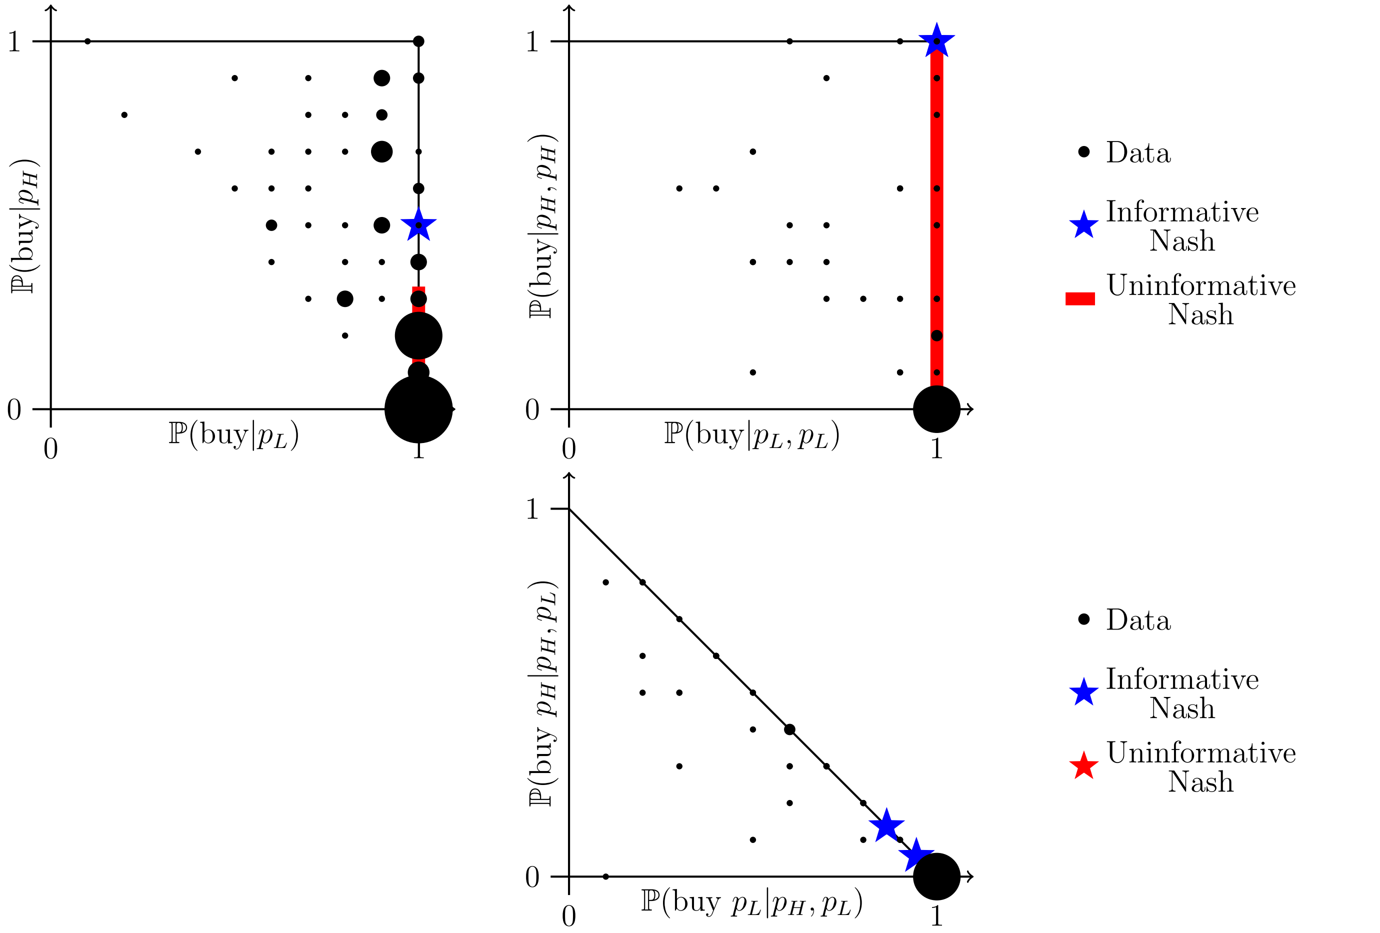

(a) $n = 1$

(b) $n = 2$

Figure 24: Empirical Buyer Strategies

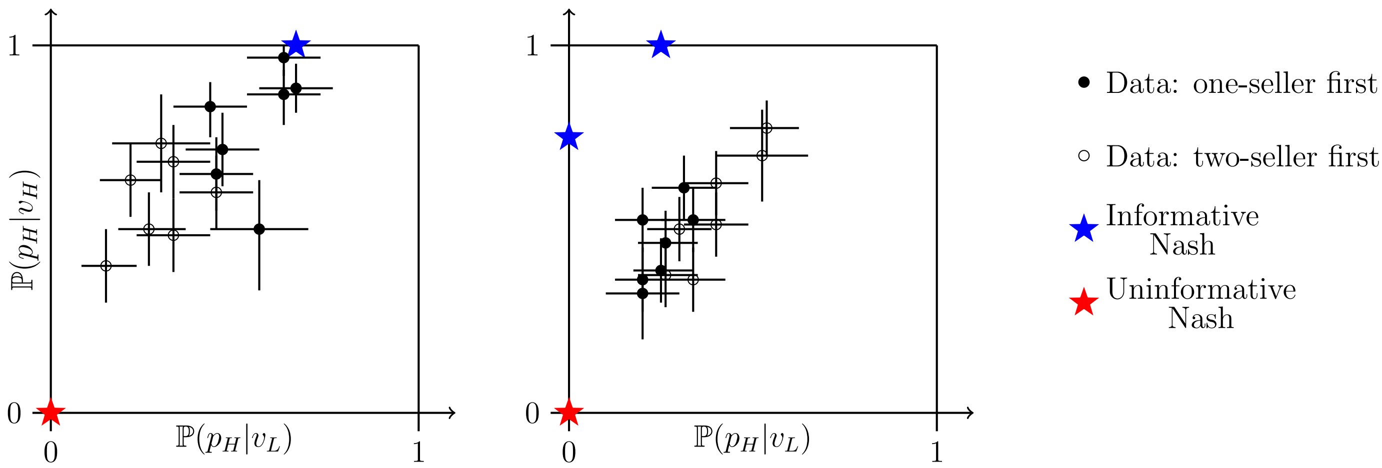

(a) $n = 1$

(b) $n = 2$

Figure 25: Empirical Seller Strategies by Session

Table 3: Permutation Tests

Outcome	Sample size	Number of permutations	More extreme permutations	p-value
High-quality price decreases	14	16384	112	0.0068
Low-quality price decreases	14	16384	1109	0.0677
Average price decreases	14	16384	285	0.0174
Informativeness decreases	14	16384	864	0.0527

Chapter 3

Cursedness in Simultaneous and Sequential Voting Games

Abstract

Agents often fail to accurately infer others' private information from their behavior, formalized as "cursedness" (Eyster and Rabin, 2005). Inferring others' types is even more difficult when agents must condition on behavior that is not directly observable. In a voting game, a voter should condition their behavior on the information set where their vote is pivotal, even if they do not yet know whether they will turn out to be pivotal. I examine voting game experiments conducted by Anderson et al. (2022) to compare simultaneous and sequential voting outcomes. I find that agents are much more cursed when they must infer others' types based on the hypothetical event in which they are pivotal, and less cursed when they see from others' behavior that they are in fact likely to be pivotal. This result suggests that the primary difficulty agents face in voting games, common-value auctions, and adverse selection environments is *not* misunderstanding the motives of other players, but rather failing to realize that their own actions only matter in "pivotal" scenarios.

1 Introduction

People often fail to infer others' private information from their behavior. This type of inconsistent belief is formalized as cursedness in Eyster and Rabin (2005). Cursedness can exist in any game of private information; Eyster and Rabin discuss its role in common-value auctions, adverse selection, and voting games.

In each of these contexts, players need to estimate the value of something that depends on other players' private information. In a common-value auction, bidders estimate the value of the prize, which depends on each bidder's private signal. In an adverse selection situation, a buyer must estimate the value of an object which is known privately to the seller. In a voting game of information aggregation, voters estimate the value of passing a reform which depends on each voter's private information.

In each case, others' behavior could reveal otherwise unknown information to players. So cursedness—a failure to infer from others' behavior—is one way players might have trouble. But in many cases, these games present an additional hurdle, related to cursedness but distinct.

In a standard common-value auction, a player receives the object only if they win the auction by outbidding the other participants. If they lose, they win nothing and pay nothing. So players estimating the value of the object are interested in its expected value conditional on winning the auction, which is to say, conditional on all the other players bidding below them. Bidders in a common-value auction should appear to under-bid relative to the unconditional value of the item because they receive the item only in a particular subset of the other players' behavior—specifically behavior that implies bad news about the value of the item. Bidders must internalize what it *would* mean if they won before choosing their bid. They need to know to condition their expectation of the value of the item on the information they would infer from winning—before they even know whether they will win or not. They must be able to infer from a particular subset of others' behavior before that behavior even occurs.

This can also be true in a voting game. A voter's vote matters only if they are pivotal, and that implies a particular profile of votes from the other players. In some cases, the situation where others vote so as to make your vote pivotal conveys information about other voters' signals and thus information about the value of passing the reform in the first place.

Because this situation involves an added hurdle, it may be that players who would correctly infer information from others' behavior if it were directly observed, might fail to correctly infer when they must construct the behavior of others that they need to be focusing on without seeing it directly. Eyster and Rabin note that if we were to measure cursedness empirically, we might measure higher levels of cursedness if agents must infer based on implied behavior rather than observed behavior.

The voting game experiments by Anderson et al. (2022) offer an opportunity to test this. By comparing simultaneous and sequential voting outcomes, we can see if the game is difficult because it is hard to infer from others' behavior or whether it is difficult because it is hard to construct the subset of the strategy space to focus on.

2 Experimental Setup from Anderson et al. (2022)

The voting games are set up like a jury, although that language is not used to participants. A hypothetical accused person is either guilty or innocent, with equal probability. Each of twelve voters gets a private signal of the guilt or innocence of the accused. The signals are generally, but not always correct. After observing their private signals, the participants each vote to convict or acquit.

The vote is not carried by a simple majority. In some treatments, the accused is not convicted unless all twelve voters vote to convict. In other treatments, a 10-2 majority is required to convict. Anything less results in acquittal. Thus, a voter is pivotal in the unanimous treatment only if all eleven other voters are voting to convict. The situation where a voter is pivotal is then potentially very informative about other's signals.

In the simultaneous treatment, voters cannot see what the others will do before casting their own vote. So they must vote as if they are pivotal: imagining that information set and inferring others' types based on it. In the sequential treatment, voters vote in order, so the kth voter sees the voting decisions of the $k-1$ others who have already voted. In this treatment, voters must still infer others' signals from their behavior, but less is left to the imagination. Instead of imagining that they may be pivotal, voters see more and more clearly that they are pivotal as more and more votes are cast.

The voting process is purely about aggregating information, not preferences. All voters get the same utility which depends only on the group decision. They each get \$4 for a correct verdict, they get \$2 for acquitting when the true state is guilt, and they get \$0 for convicting when the true state is innocence. Conditional on the true state, the voters' signals are i.i.d. and are correct with 75% probability.

3 Voting Game Theory

The theory of the voting games is a story about how the inability of voters to communicate their signals, combined with a sub-optimal voting rule (unanimous or a 10-2 majority), makes the first-best outcome impossible. All voters agree that the accused should be convicted if the number of signals that indicate guilt relative to innocence outweigh the relative harm of Type I error versus Type II error. So if voters could freely communicate their signals, they would all cooperate on the best outcome (in this specific case, they could all vote to convict if 7 or more of the signals indicated guilt).

But even when voters cannot communicate, the voting mechanism could potentially aggregate their information perfectly and still reach the first-best outcome of conviction if and only if 7 or more voters received a signal of guilt. To do this, the voting rule must allow all voters to reveal their signal, since not all information can be aggregated if not all information is revealed to the mechanism. And this is where the voting rules fail in this case.

Suppose unanimity is required to convict, and anything less results in acquittal. If everyone votes according to the signal they received, then conviction will only occur if everyone received signals of guilt. But this is too stringent a requirement—it is best if conviction occurs if 7 or more voters receive signals of guilt. So if everyone reveals their signals to the mechanism, the mechanism often will choose the wrong outcome. [1]

An optimal voter will correct for this failure of the voting mechanism by not truthfully revealing their signal. Sometimes, upon receiving a signal of innocence, the voter will still vote to convict. This happens because optimal voters know that their vote matters only if they are pivotal, and so they assume they will be pivotal when deciding how to vote (Fedderson and Pesendorfer, 1998).

Suppose again that unanimity is required to convict, and suppose that voters are naïvely

[1] The optimal voting rule in this case is a simple majority, with ties resulting in acquittal. This allows the correct outcome to be selected when all voters tell the truth about their signals to the mechanism. But in general, the optimal voting rule will depend on the number of voters, the relative harm of type I vs. type II error, and the strength of each signal.

voting according to the signal they receive. A voter should realize that, unless all others vote to convict, the accused will be acquitted regardless of how they themselves vote; their vote matters only if everyone else is voting to convict. But if everyone else is voting to convict, that means they all received signals of guilt. This is such overwhelming evidence of guilt that the voter is better off voting to convict, even if they themselves received a signal of innocence.

Thus, voters voting according to their signals is not a Nash equilibrium. An equilibrium requires that the situation of being pivotal convey not so much information that it causes voters to ignore their own signal completely. In a symmetric, informative Nash equilibrium, voters will vote to convict upon receiving a signal of guilt and mix upon receiving a signal of innocence.[2]

It might seem that switching from simultaneous to sequential voting would completely change the Nash equilibria since it changes the information available to voters. In fact, Dekel and Piccione (2000) show that the set of Nash equilibria is almost unchanged. In particular, the simultaneous voting equilibria are all still equilibria under sequential voting. Under the unanimous voting rule, the equilibria are exactly the same.

This happens because voters should vote as if they are pivotal. Even if they don't turn out to be pivotal, their vote won't matter, so they should assume they will be pivotal when deciding how to vote. Under sequential voting, voters are given information that indicates they are more and more likely to be pivotal. If the path of votes serves only to update voters' beliefs that they will be pivotal, it will not affect their decision. They are already conditioning on the state where they are pivotal.[3]

[2]The same intuition holds in non-symmetric equilibria: being pivotal cannot convey too much information. A table of all informative equilibria for the unanimous voting rule is given in the appendix.

[3]Under the weighted majority voting rule, the path of votes can actually convey slightly more that just how close to pivotal voters are becoming. It can reveal to voters how they are pivotal. Multiple different strings of past votes will lead to a voter being pivotal and sequential voting allows voters to see which one is occurring. This could lead to additional equilibria that exist in the sequential treatment.

4 The Data (Anderson et al., 2022)

A total of 400 games were played, 200 with simultaneous voting and 200 with sequential voting. In each case, 100 games used the unanimous voting rule and 100 used the 10-2 majority voting rule. In the simultaneous treatment, players generally voted their signal but were noisy. Under both the unanimous and weighted majority rules, players voted their signal about 80% of the time. They do not appear to correct at all for the flawed voting mechanism. Even though the mechanism favors acquittal, players follow their signal with about the same probability whether it is a signal of guilt or innocence. As a result, nearly every game in the simultaneous treatment results in acquittal.

Table 1: **Simultaneous Voters are not Strategic**

Voting Rule	$\mathbb{P}$**(vote convict**$\vert \boldsymbol{i}$**)**	$\mathbb{P}$**(vote convict**$\vert \boldsymbol{g}$**)**
Unanimous:	0.19	0.78
Majority (10-2):	0.15	0.80

Notes: Table reports empirical choice probabilities in the simultaneous voting games. i represents a signal of innocence and g represents a signal of guilt. Voters seem to completely ignore the skewed voting rule and simply vote according to their signal about 80% of the time.

In the sequential treatment, people vote in order, and see the sequence of votes that have already been cast. If a decision is reached before all twelve voters have voted, the game ends. For instance, if the first voter votes to acquit under the unanimous voting rule, the accused will be acquitted regardless of how the others vote, so the game ends and the others do not vote.

Thus, if the game has not already ended, players who vote later in the sequence are more likely to be pivotal and can see this from the path of votes. Since the voting rules are skewed to make conviction more difficult, a voter who votes late in the sequence will see many votes to convict and only one or two votes to acquit in the sequence of votes already cast.

Voters respond to this by generally voting according to their signal early in the sequence, but leaning more and more towards conviction as they see more others voting to convict. The eleventh and twelfth voters, if their votes still matter, almost always vote to convict regardless of their private signal. On average, individual voters in the sequential treatment fight the skewed voting rules by voting to convict much more than to acquit.

Table 2: **Sequential Voters are Increasingly Strategic**

	Unanimous		**Majority (10-2)**	
Voting Order	$\mathbb{P}(\boldsymbol{c}\vert\boldsymbol{i})$	$\mathbb{P}(\boldsymbol{c}\vert\boldsymbol{g})$	$\mathbb{P}(\boldsymbol{c}\vert\boldsymbol{i})$	$\mathbb{P}(\boldsymbol{c}\vert\boldsymbol{g})$
1	0.16	0.93	0.00	0.90
2	0.41	1.00	0.02	0.98
3	0.67	0.86	0.02	0.93
4	0.73	1.00	0.18	0.87
5	0.60	1.00	0.38	0.89
6	0.83	1.00	0.57	0.94
7	1.00	1.00	0.62	0.91
8	0.64	1.00	0.50	0.93
9	0.88	1.00	0.67	1.00
10	0.70	1.00	0.89	0.96
11	0.88	1.00	0.77	1.00
12	0.89	1.00	1.00	0.92

Notes: Table reports empirical choice probabilities in the sequential voting games, broken down by position in the voting order. i represents a signal of innocence and g represents a signal of guilt. Voters seem to vote to convict more and more the later they are in the voting order (and thus the more convict votes they see already cast).

The sequential treatment makes the game easier for participants by explicitly showing others' behavior. In the following sections, I will try to show that the difference between the sequential and simultaneous data is a result of differences in the consistency of voters' beliefs, measured by cursedness.

5 Explaining the Data with Noise Alone

First, I'll show what it looks like to try to explain the data with just noise. Here I assume that voters have consistent beliefs, but their actions are not always optimal. I implement this with the logit quantal response equilibrium (McKelvey and Palfrey, 1995). A QRE works by adding noise to players' perceptions of their payoffs. This means a player's strategy will be distributed around their best response. In the logit QRE, the payoff errors are Extreme

Value Type I. The probability that a logit quantal response agent will choose action x from strategy space X is

$$\sigma(x) = \frac{e^{\lambda u(x)}}{\sum_{i \in X} e^{\lambda u(i)}} \tag{1}$$

Here λ is the inverse of the scale parameter for the logit errors. It represents the precision of the agents. If λ is close to 0, the errors have high variance and the noise overwhelms the underlying payoffs. In this case, an agent will choose actions from X uniformly randomly. As $\lambda \to \infty$, the variance of the errors goes to zero, and agents converge to choosing their best response with probability 1. When this happens, the QRE converges to a Nash equilibrium.

In our case, an agent receives a signal s (either a signal of innocence or guilt). The probability that a voter with signal s will vote to convict is

$$\sigma_s = \frac{e^{\lambda \mathbb{E} u(\text{vote convict})}}{e^{\lambda \mathbb{E} u(\text{vote convict})} + e^{\lambda \mathbb{E} u(\text{vote acquit})}} \tag{2}$$

If $\lambda = 0$, agents will vote to convict with probability $\frac{1}{2}$. As λ increases, they will approach the Nash equilibrium of voting to convict with probability 1 if they received a signal of guilt and some probability less than 1 if they received a signal of innocence. Figure 1 shows the symmetric, informative QRE correspondence for the unanimous voting rule. There are two things that are instructive here.

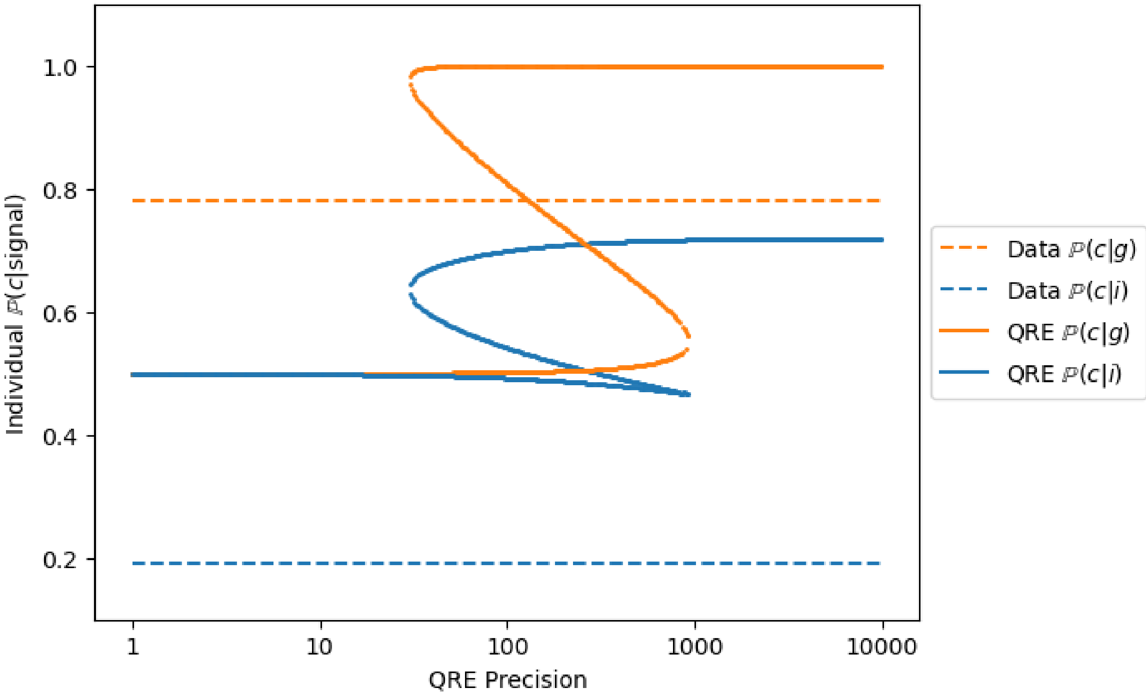

Figure 1: Noise Cannot Rationalize Voting to Acquit

Notes: Figure shows the quantal response equilibrium correspondence under the unanimous voting rule. c represents voting to convict, while i represents a signal of innocence and g represents a signal of guilt. Importantly, at no precision does the predicted probability of voting to convict fall below about 48%, far above the empirical probability of 19% with a signal of innocence.

First, the QRE correspondence does not converge to the Nash equilibrium until precision is extremely high. In other words, the noise overwhelms the true payoff difference even when the errors have low variance. The expected payoff of voting to convict is almost the same as the expected payoff of voting to acquit. This is because it's very unlikely that an individual vote will matter. The chance of a voter being pivotal is about 17% at the Nash equilibrium if the voter has a signal of innocence, but if precision is lower, the probability decreases towards 0.0005.

A QRE close to the Nash equilibrium appears around $\lambda = 56$, but it does not become unique until about $\lambda = 1216$. Empirically, the probability that a voter is pivotal is 0.0016 if they have a signal of innocence. If the voter's choice doesn't matter over 99% of the time, the expected payoff difference between voting to convict and to acquit is very small. It comes out to a difference of about 1 cent (if they have a signal of guilt) or 0.3 cents (if they have a signal of innocence). When we see agents systematically choosing convict or acquit based on their signal, the QRE would have to conclude that they are unrealistically precise. The alternative is that voters' mistakes are more than just noisy behavior. Perhaps voters do not fully appreciate that for their vote to matter, all eleven other voters have to choose to convict.

The second thing to notice about the graph is that, in any QRE at any precision, voters do not vote to convict less than about 48% of the time when they have a signal of innocence. This is because, roughly speaking, QRE agents are on a spectrum between uniform randomization and Nash equilibrium. In the Nash equilibrium, agents correct for the flawed voting rule by voting to convict sometimes, even when they have a private signal of innocence. In the unanimous treatment, they should be voting to convict 72% of the time when they have an innocent signal. On the other hand, if agents are imprecise, they will uniformly randomize. In this case, they will vote to convict 50% of the time, regardless of what signal they have.

The QRE says that the only way for agents to be suboptimal is for them to be noisy. If I analyze the game with just noise as an explanation for suboptimal behavior, there will be no way to explain persistent, specific suboptimal choices. But this is what we see in the data. When voters vote simultaneously under the unanimous voting rule, those who receive a signal of innocence vote to convict 19% of the time. This is far below the 50% probability from uniform randomization and even farther from the 72% probability of the Nash. Voters are being fairly precise in their inaccuracy. In this case, voter are behaving worse than randomly. They are systematically choosing a strategy that gives them a lower payoff than if they had just flipped a coin to decide what action to choose. This systematic homing in on a particular suboptimal strategy cannot be explained by noise.[4]

[4]On the other hand, when a voter receives a signal of guilt in this treatment, they vote to convict 78% of the time. The QRE *can* explain this. A totally imprecise voter would randomize and vote to convict

Simultaneous voting under the weighted majority voting rule cannot be explained by noise for the same reason, although it is less obvious simply because it is easier to be pivotal. People behave differently under the sequential treatment, and it seems that QRE may be able to explain these treatments without recourse to inconsistent beliefs. The biggest mistakes are specific to simultaneous voting.

6 Adding Belief Inconsistency

The next step is to add inconsistency into beliefs and see how it affects simultaneous vs. sequential voting. The difficult part of this game is realizing how unlikely it is to be pivotal—it only occurs in select information sets—and then leveraging that information into a posterior probability of guilt. So it is natural to introduce inconsistency into beliefs using cursedness. Cursedness is simply the failure to leverage the informational content of others' behavior.

In fact, in this case I cannot use canonical cursedness exactly. Instead, I introduce the failure of inference by assuming agents underestimate the number of other players in the game. Consider again the case of simultaneous voting under the unanimous voting rule. Under this rule, a voter is only pivotal if all others vote to convict. If (just for the sake of argument) voters are voting according to their signals, a voter is only pivotal if all others have received signals of guilt. This is much more conclusive evidence if there are 11 other voters than if there are only 2 other voters.

Because this is not the exact formulation of cursedness used by Eyster and Rabin, I call it *strategic unawareness*–but intuitively it is doing the exact same job as cursedness. Strategic unawareness leads voters to ignore the weight of evidence conveyed by being pivotal, just like cursedness. In fact, applying cursedness is generally mathematically identical to applying strategic unawareness. But in a model which also includes noise from quantal response, the

with 50% probability, but very precise voter will vote to convict with probability 1. A probability of 78% is rationalized by an intermediate precision.

two approaches can yield slightly different results.

Under strategic unawareness, voters do not realize how unlikely it is to be pivotal. They think their vote is likely to matter in lots of situations, and precisely because their vote matters in lots of situations, being pivotal does not convey much information. Under canonical cursedness, voters realize that their chance of being pivotal is tiny, and yet they still discount the informational content of that situation. This has the potential to interact perniciously with the noise from quantal response.

If voters are quantal response agents who are also cursed, it will be hard to understand their level of noisiness. On the one hand, it appears that they occasionally make mistakes (such as failing to vote convict when receiving a signal of guilt), but generally do not. They seem reasonably precise, but not incredibly so. On the other hand, cursed voters understand that the probability of their vote mattering is tiny. So the expected payoff difference between voting convict and voting acquit is similarly tiny, especially if they have a signal of innocence. For voters to even care about their vote when it's almost completely irrelevant demonstrates massive precision. And this is hard to reconcile intuitively with how voters seem to act. I tend to think that voters fail to realize in the first place that their vote won't matter unless others all vote to convict, and *as a result* they do not take into account the information that such a situation would convey. Thus, I fit strategic unawareness rather than canonical cursedness.

Both ways of implementing belief inconsistency push voters toward voting according to their signal. As they discount the information from others, their own signal becomes more and more important to their decision. Like cursedness, I have put strategic unawareness on a scale from 0 to 1, where 0 is fully consistent beliefs and 1 is completely ignoring the information from other players.

Figure 2 shows the choice probabilities under the unanimous voting rule as strategic unawareness goes from 0 to 1. Agents start at the Nash probabilities and move to eventually voting according to their signals once strategic unawareness is high. There are other equilibria besides this symmetric one, but they all involve players as a whole moving from the Nash towards voting their signal.

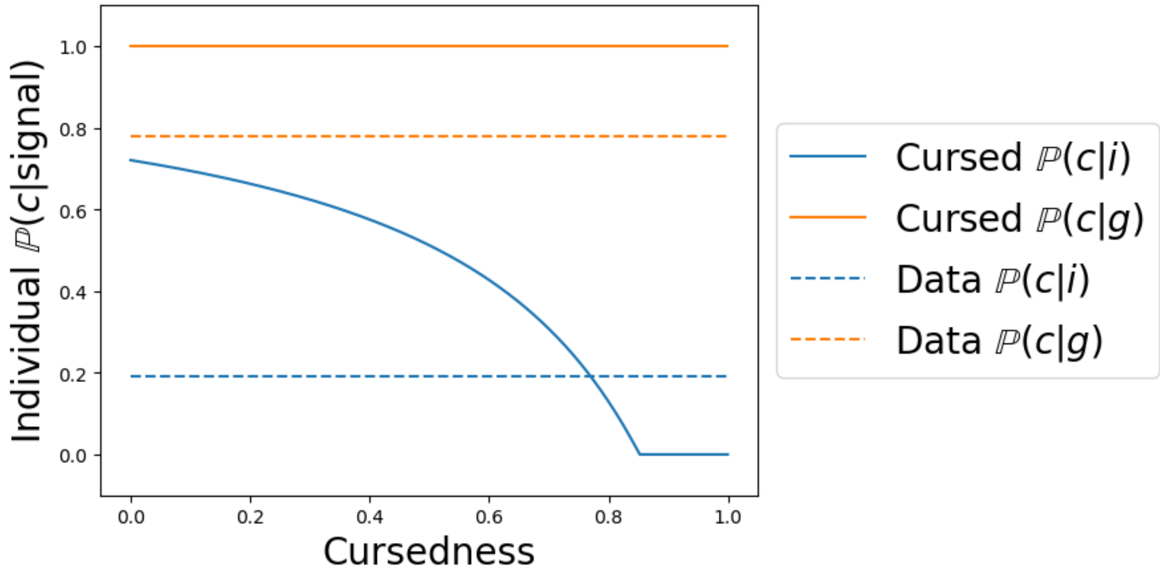

Figure 2: Cursedness Explains Non-Strategic Voting

Notes: Figure shows the cursed equilibrium under the unanimous voting rule as the level of cursedness varies. (Since this graph does not involve noise, cursedness and strategic unawareness are identical.) c represents voting to convict, while i represents a signal of innocence and g represents a signal of guilt. As individuals are more cursed, they ignore the informational weight of being pivotal, and thus are more and more likely to vote according to their private signal.

Figure 2 ignores noise, which is why the probability to vote convict when receiving a signal of guilt is always 1. Since people do not always vote to convict when they have a guilty signal, the likelihood of the data is 0 given this model without noise. We need noise both to explain the significant variance of behavior and also to have a well-behaved likelihood function to allow maximum likelihood estimation of strategic unawareness.

Adding strategic unawareness into the QRE for simultaneous voting yields a strategic unawareness of 0.89 and a precision of 13.65. The low value of precision indicates that the game is complicated in general, and so players have a lot of variance around their best response. But the high strategic unawareness indicates that players are specifically unable to infer information from being pivotal. They ignore about 90% of the weight of guilt that would result from being pivotal.

Adding strategic unawareness into the QRE for sequential voting tells a very different story. First of all, there are many paths of votes that never occur, and thus we do not know how agents would have responded to them. This makes the construction of beliefs difficult, so that not every node for which we have data can be used in the likelihood calculation.

But if we restrict our attention to those nodes where we can calculate the likelihood, we measure a precision of 13.15 and a strategic unawareness of -0.04–not significantly different from 0. Suddenly, while players are still equally noisy, their problems can no longer be traced to inconsistency of beliefs, or at least not a failure to infer information from others' behavior.

Table 3: **Precision and Strategic Awareness Joint Estimates**

Timing	Precision λ	Strategic Unawareness χ
Simultaneous:	13.65	0.89
Sequential:	13.15	-0.04

7 Conclusion

These voting games are difficult for players to think through. Players are both noisy in general–unresponsive to small differences in expected payoffs–and also particularly bad at responding to the weight of evidence implied by being pivotal to the group decision. However, when voting is conducted sequentially, so that subjects can see the decisions of those who have already voted, players do a much better job internalizing what it means to be pivotal.

In both simultaneous and sequential voting, players must infer others' private information from their behavior. The only difference is that, when voting is simultaneous, players must realize that their vote will rarely sway the group decision, and yet their decision only matters when they are in a position to swing the group decision. Thus, simultaneous voting requires agents to imagine what behavior from other voters would lead to their own vote being pivotal, whereas sequential voting allows voters to see the string of others' votes that make their own vote pivotal.

The huge difference in estimated cursedness (technically strategic unawareness) between simultaneous and sequential voting shows that inference from others' behavior to others' private information is not the problem. Subjects are fairly good at guessing others' private information from their behavior when they observe that behavior directly. Instead, strategic voting is difficult because agents fail to understand that their vote only matters when they are pivotal.

Appendices

A Derivation of equilibria

The voting games are structured like a jury where a supermajority is required to convict, and anything less results in acquittal. There is first of all a true state, innocent (I) or guilty (G), unknown to the voters. Each state is equally likely (although it is easy to generalize to arbitrary prior probabilities). Each of n voters then receives a private signal of innocence (i) or guilt (g). The signals are i.i.d. conditional on the true state, and are correct with probability p. The voters choose to vote to acquit (a) or convict (c), and if at least k voters have voted to convict, the group decision will be conviction (C). If less than k vote to convict, the group decision will be acquittal (A).

Each voter has the same preferences, so the voting is purely an information aggregation mechanism. A good outcome (acquittal when the true state is innocent or conviction when the true state is guilty) yields a payoff of zero. A type I error (conviction when innocent) yields a payoff of $-q$ and a type II error (acquittal when guilty) yields a payoff of $-(1-q)$.

Consider an agent participating in the simultaneous voting mechanism. There are two aspects of this game that make it difficult to reach the optimal strategy. First, the agent must impute the behavior of others that would make his own vote matter. An individual voter's choice will only affect the group decision if $k-1$ other voters are choosing to convict and $n-k$ others are choosing to acquit: otherwise, the outcome is already decided and the agent's choice does not matter. A rational agent will restrict their attention to this information set when deciding how to vote.

Second, the agent must calculate the value of conviction versus acquittal conditional on this information set. If exactly $k-1$ others are voting to convict and $n-k$ are voting to acquit, and the voters are taking their private information into account in their voting decision, the agent should be able to infer something about the others' private information.

A rational agent will realize that the fact of being pivotal is news about the guilt or innocence of the accused and take that news into account in their individual voting decision.

Mathematically, an agent should choose to vote c whenever

$$\mathbb{E}[u(c)–u(a)] > 0$$

Since the probability of being pivotal does not depend on the agent's own vote, this expectation can be decomposed into

$$\mathbb{P}(\text{not pivotal})\mathbb{E}[u(c) - u(a)|\text{not pivotal}] + \mathbb{P}(\text{pivotal})\mathbb{E}[u(c) - u(a)|\text{pivotal}]$$

In the first case, if the agent is not pivotal then their vote does not change the group decision. Since their utility depends only on the group decision, $u(c) = u(a)$ when the agent is not pivotal. When the agent is pivotal, the group decision is determined by their vote, so $u(c) = u(C)$ and $u(a) = u(A)$: the utility of voting convict is the utility of reaching a group decision of conviction. As long as there is a nonzero chance that the agent's vote will be pivotal, they should vote to convict when

$$\mathbb{E}[u(C) - u(A)|k - 1 \text{ voted } c,\ n - k \text{ voted } a] > 0$$

The agent should simply assume they are pivotal when deciding how to vote.

Here I am considering cases where k is large relative to n: a supermajority are required to reach conviction. I am also looking at cases where q is small enough that a supermajority of guilty signals would make conviction preferable. (If q is too large, even $k - 1$ out of n guilty signals will not make conviction preferable since the harm of accidentally convicting the innocent is so much worse than accidentally acquitting the guilty.)

The fact that rational agents take into account the information conveyed by being pivotal will bound how much information can be aggregated in equilibrium. If players are generally voting to convict when they have a signal of guilt and voting to acquit when they have a signal of innocence, then being pivotal conveys massive information. A voter is only pivotal when $k-1$ others are voting to convict. If voters are generally following their signals, this means a supermajority must have received signals of guilt. In this case, a voter might be better off ignoring their own signal since so many others would have to have received signals of guilt if the voter were to be pivotal. So it cannot be an equilibrium for voters to generally vote according to their signal.

In a symmetric equilibrium, found by Federson and Pesendorfer (1998), agents vote to convict when they receive a signal of guilt, but randomize when they receive a signal of innocence. The randomization probability makes the fact of being pivotal convey just enough evidence of guilt so that agents are indifferent as to whether to follow their signal or just vote to convict. The exact probability of voting to convict when observing a signal of innocence, σ_0, is given by

$$\begin{aligned}
\mathbb{E}[u(C)|i,\ \text{pivotal}] &= \mathbb{E}[u(A)|i,\ \text{pivotal}] \\
qp\mathbb{P}(\text{pivotal}|I) &= (1-q)(1-p)\mathbb{P}(\text{pivotal}|G) \\
qp\gamma_0^{k-1}(1-\gamma_0)^{n-k} &= (1-q)(1-p)\gamma_1^{k-1}(1-\gamma_1)^{n-k} \\
\text{where } \gamma_0 = 1-p+p\sigma_0,&\ \ \gamma_1 = p+(1-p)\sigma_0
\end{aligned}$$

This is an equilibrium because being pivotal conveys roughly enough information to make agents indifferent about whether to follow their signal. It does not matter whether that information is conveyed by slightly informative decisions from every voter or very informative decisions from just a few voters. So there are a host of asymmetric equilibria where some agents are more likely to vote according to their signal than others.

In particular, there is a pure strategy asymmetric equilibrium where m agents vote according to their signals and the rest always vote to convict. m is defined as the integer(s) satisfying

$$2k - n - 2 - \frac{ln\left(\frac{q}{1-q}\right)}{ln\left(\frac{p}{1-p}\right)} \leq m \leq 2k - n - 1 - \frac{ln\left(\frac{q}{1-q}\right)}{ln\left(\frac{p}{1-p}\right)}$$

Now it is clear why the sequential voting game does not change these equilibria. In the simultaneous game, agents are already assuming they will be pivotal when casting their vote. The difference in the sequential game is that agents can see directly that they are more likely to turn out to be pivotal. Since they were assuming they would be pivotal anyway, this information should not change their behavior.

When $k < n$, there are other equilibria that arise in the sequential game that are not equilibria in the simultaneous game. In these equilibria, earlier voters vote differently because they think later voters will care not just about how many convict and acquit votes have occurred so far, but also about the specific order of the votes. This belief could become self-fulfilling and lead to an equilibrium that is specific to sequential play. Dekel and Piccione do not solve for these additional equilibria and for now I do not consider them, but I hope to analyze them in the future.

B Optimal voting rule

For the parameters used by Anderson et. al, the first-best outcome is for conviction to occur whenever there are 7 or more signals of guilt. This is because, if k of the n signals indicate guilt, the probability of guilt is

$$\begin{aligned}
\mathbb{P}(G|k) &= \frac{\mathbb{P}(k|G)\mathbb{P}(G)}{\mathbb{P}(k|G)\mathbb{P}(G) + \mathbb{P}(k|I)\mathbb{P}(I)} \\
&= \frac{\binom{n}{k}p^k(1-p)^{n-k}\left(\frac{1}{2}\right)}{\binom{n}{k}p^k(1-p)^{n-k}\left(\frac{1}{2}\right) + \binom{n}{k}p^{n-k}(1-p)^k\left(\frac{1}{2}\right)} \\
&= \frac{p^k(1-p)^{n-k}}{p^k(1-p)^{n-k} + p^{n-k}(1-p)^k}
\end{aligned}$$

and similarly

$$\mathbb{P}(I|k) = \frac{p^{n-k}(1-p)^k}{p^k(1-p)^{n-k} + p^{n-k}(1-p)^k}$$

All agents want conviction to occur when

$$\begin{aligned}
u(C|k) &> u(A|k) \\
\implies -q\mathbb{P}(I|k) &> -(1-q)\mathbb{P}(G|k) \\
\implies -qp^{n-k}(1-p)^k &> -(1-q)p^k(1-p)^{n-k} \\
\implies \left(\frac{q}{1-q}\right)\left(\frac{p}{1-p}\right)^{n-k} &< \left(\frac{p}{1-p}\right)^k \\
\implies \left(\frac{q}{1-q}\right) &< \left(\frac{p}{1-p}\right)^{2k-n} \\
\implies ln\left(\frac{q}{1-q}\right) &< (2k-n)ln\left(\frac{p}{1-p}\right) \\
\implies k &> \frac{n}{2} + \frac{ln\left(\frac{q}{1-q}\right)}{2ln\left(\frac{p}{1-p}\right)}
\end{aligned}$$

Since $q < p$, the fraction on the right is smaller than 1, so conviction will be best for any $k > n/2$, which in our case implies $k \geq 7$ should lead to conviction.

C Additional Figures

Table 4 shows all the informative equilibria under the unanimous voting rule. Equilibria are the same regardless of whether voting is simultaneous or sequential. In general, pivotality cannot convey too much information in equilibrium, or individuals will choose to ignore their private information and follow the weight of evidence implied by pivotality. Equilibria can differ in *how* pivotality conveys information: whether through many voters who are slightly informative or a few voters who are very informative.

The table shows that asymmetric equilibria are made up of some number of voters (m_i)

who vote according to their signal and are thus perfectly informative, some number of voters (m_σ) who vote to convict when observing a signal of guilt and mix when observing a signal of innocence and are thus partially informative, and some number of voters (m_1) who always vote to convict regardless of their signal, and are thus uninformative.

The common thread is that there cannot be too much information conveyed. There cannot be too many informative voters (m_i can never be above two), and the more partially informative voters there are (higher m_σ) the less informative each one can be (lower σ). Partially informative voters must be indifferent in order to mix when observing a signal of innocence, so for those voters the posterior probability of guilt that comes through being pivotal must be exactly equal to the relative harm of type I versus type II error. This is why, in every equilibrium, $\mathbb{P}(G|\text{pivotal}, i) = 2/3$ for these voters.

Mathematically, suppose a voter is pivotal when there are m_i other informative voters and m_σ other partially informative voters, and suppose the voter has a signal of innocence. The probability that the accused is guilty is then

$$\begin{aligned}\mathbb{P}(G|\text{pivotal}, i) &= \frac{\mathbb{P}(\text{pivotal}|G, i)\mathbb{P}(G|i)}{\mathbb{P}(\text{pivotal}|i)} \\ &= \frac{\mathbb{P}(\text{pivotal}|G)\mathbb{P}(G|i)}{\mathbb{P}(\text{pivotal}|G)\mathbb{P}(G|i) + \mathbb{P}(\text{pivotal}|I)\mathbb{P}(I|i)} \\ &= \frac{\mathbb{P}(\text{pivotal}|G)(1-p)}{\mathbb{P}(\text{pivotal}|G)(1-p) + \mathbb{P}(\text{pivotal}|I)p} \\ &= \frac{p^{m_i}(p+(1-p)\sigma)^{m_\sigma}(1-p)}{p^{m_i}(p+(1-p)\sigma)^{m_\sigma}(1-p) + (1-p)^{m_i}(1-p+p\sigma)^{m_\sigma}p}\end{aligned}$$

Note that there are also uninformative equilibria, where the outcome is the same regardless of everyone's signals. For instance, if k out of n votes are required to convict, then it is an equilibrium for $n-k+2$ or more voters to always vote acquit, regardless of their signal. The outcome is that the accused will always be acquitted, and no one voter can unilaterally change the outcome: there will still be at least $n-k+1$ others voting to acquit. Since no one's vote matters, any behavior is privately optimal, so always voting to acquit is a best

Table 4: **Non-trivial Nash Equilibria:**

Player Strategies				$\mathbb{P}(G\|\text{pivotal}, i)$		
m_i	m_σ	m_1	σ	**informative**	**mixed**	**uninformative**
0	3	9	0.09		0.67	0.83
0	4	8	0.27		0.67	0.78
0	5	7	0.39		0.67	0.76
0	6	6	0.48		0.67	0.74
0	7	5	0.54		0.67	0.73
0	8	4	0.59		0.67	0.72
0	9	3	0.64		0.67	0.71
0	10	2	0.67		0.67	0.71
0	11	1	0.70		0.67	0.71
0	12	0	0.72		0.67	
1	2	9	0.20	0.57	0.67	0.80
1	3	8	0.49	0.49	0.67	0.74
1	4	7	0.63	0.46	0.67	0.72
1	5	6	0.71	0.44	0.67	0.70
1	6	5	0.76	0.43	0.67	0.70
1	7	4	0.79	0.43	0.67	0.69
1	8	3	0.82	0.42	0.67	0.69
1	9	2	0.84	0.42	0.67	0.69
1	10	1	0.86	0.42	0.67	0.68
1	11	0	0.87	0.42	0.67	
2	0	10		0.50		0.75

Notes: This table shows the informative equilibria in the unanimous voting game (both simultaneous and sequential have the same equilibria in this case). In each equilibrium, the twelve voters are made up of m_i voters who vote according to their signal, m_1 voters who always vote to convict, and m_σ voters who vote to convict when they have a signal of guilt, and vote convict with probability σ when they have a signal of innocence.

response to at least $n - k + 1$ others always voting to acquit. The same situation is true if there are more than $k + 1$ players who always vote to convict, regardless of their signal. No one will be able to unilaterally change the outcome, so these situations will all be Nash equilibria.

Conclusion

In each of the strategic situtations in this work, Nash agents could have reached an uninformative equilibrium where private information does not influence actions. Instead, agents almost always reached an equilibrium where they behave strategically, choosing based on their private information but at the same time not being *perfectly* informative. In the pricing games, perfect informativeness does not occur because above a certain level, more informative pricing would not serve the overarching goal of maximizing seller profits. In the voting games, informativeness above a certain level would not serve the overarching goal of convicting those most likely to be guilty and acquitting those most likely to be innocent.

At the same time, agents are plagued by different types of behavioral bias. The most obvious of these is noisy decision-making. A small amount of noise can quickly erode informativeness, since agents will occasionally make choices that do not correspond with their private information. This can decrease informativeness, even if choice probabilities unconditional on private information do not change much. This is especially true with binary decisions, where every 'mistake' with one private signal looks like an intentional choice based on a different private signal. Noise erodes informativeness directly, and also makes it difficult for agents to deduce others' private information.

But even in the absence of noise, agents have trouble deducing others' information from their behavior, through cursedness/strategic unawareness. Sometimes, as in the case of the voting game outcomes, this failure of inference leads to suboptimal behavior beyond what could ever result from noise. Other times, agents behave noisily even when they are not required to infer anything from agent behavior.

Both behavioral biases can have complex equilibrium effects, as agents expect and respond to their fellow players' noise and failure of inference. But in general, noise erodes informativeness, while cursedness can sometimes lead to increased informativeness. This happens in both the pricing games and the voting games. In the pricing games, cursedness among buyers means prices must be more informative before cursed buyers realize prices are informative enough to make them indifferent to buying the product. In the voting games, cursedness among voters makes them underestimate the amount of information already conveyed by being pivotal, and thus allows them to act out their own private information without realizing they are dooming the group decision.

Bibliography

[1] L. R. Anderson, C. A. Holt, K. K. Sieberg, and B. A. Freeborn. An experimental study of strategic voting and accuracy of verdicts with sequential and simultaneous voting. *Games*, 13(2):26, 2022.

[2] Kyle Bagwell and Michael H Riordan. High and declining prices signal product quality. *American Economic Review*, pages 224–239, 1991.

[3] Steven T Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262, 1994.

[4] Helmut Bester and Klaus Ritzberger. Strategic pricing, signalling, and costly information acquisition. *International Journal of Industrial Organization*, 19(9):1347–1361, 2001.

[5] W Bradford Cornell. Price as a quality signal: Some additional experimental results. *Economic Inquiry*, 16(2):302, 1978.

[6] E. Dekel and M. Piccione. Sequential voting procedures in symmetric binary elections. *Journal of Political Economy*, 108(1):34–55, 2000.

[7] E. Eyster and M. Rabin. Cursed equilibrium. *Econometrica*, 73(5):1623–1672, 2005.

[8] Eugene F Fama. Efficient capital markets. *Journal of Finance*, 25(2):383–417, 1970.

[9] T. Feddersen and W. Pesendorfer. Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. *American Political Science Review*, 92(1):23–35, 1998.

[10] Sanford J Grossman and Joseph E Stiglitz. On the impossibility of informationally efficient markets. *The American economic review*, 70(3):393–408, 1980.

[11] S. Guarnaschelli, R. D. McKelvey, and T. R. Palfrey. An experimental study of jury decision rules. *American Political Science Review*, 94(2):407–423, 2000.

[12] Sander Heinsalu. Competitive pricing despite search costs when lower price signals quality. *Economic Theory*, 71(1):317–339, 2021.

[13] C. A. Holt, R. Sahu, and A. M. Smith. An experimental analysis of risk effects in attacker-defender games. *Southern Economic Journal*, 89(1):185–215, 2022.

[14] Maarten CW Janssen and Santanu Roy. Signaling quality through prices in an oligopoly. *Games and Economic Behavior*, 68(1):192–207, 2010.

[15] Harold J Leavitt. A note on some experimental findings about the meanings of price. *The Journal of Business*, 27(3):205–210, 1954.

[16] Giovanni Mastrobuoni, Franco Peracchi, and Aleksey Tetenov. Price as a signal of product quality: Some experimental evidence. *Journal of Wine Economics*, 9(2):135–152, 2014.

[17] J Douglas McConnell. An experimental examination of the price-quality relationship. *The Journal of Business*, 41(4):439–444, 1968.

[18] Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.

[19] Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for extensive form games. *Experimental economics*, 1(1):9–41, 1998.

[20] Jerry C Olson. Price as an informational cue: Effects on product evaluations. *Consumer and Industrial Buying Behavior*, 1977.

[21] Robert A Peterson. The price-perceived quality relationship: Experimental evidence. *Journal of Marketing Research*, 7(4):525–528, 1970.

[22] Paul A Samuelson. Industrial management review.“proof that properly anticipated prices fluctuate randomly”, 1965.

[23] Jean Tirole. *The theory of industrial organization.* MIT press, 1988.

[24] Theodore L. Turocy. A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence. *Games and Economic Behavior*, 51(2):243–263, 2005. Special Issue in Honor of Richard D. McKelvey.

[25] Theodore L Turocy. Computing sequential equilibria using agent quantal response equilibria. *Economic Theory*, 42:255–269, 2010.

[26] VK Venkataraman. The price-quality relationship in an experimental setting. *Journal of Advertising Research*, 1981.